Dell PowerVault MD Series Storage Arrays Administrator's Guide

Notes, Cautions, and Warnings

NOTE: A NOTE indicates important information that helps you make better use of your computer.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2012 Dell Inc.

Trademarks used in this text: Dell™, the Dell logo, Dell Boomi™, Dell Precision™ , OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation.

2012 - 11

Rev. A00

Contents

Notes, Cautions, and Warnings....2

1 Introduction....15
- Dell PowerVault Modular Disk Storage Manager....15
- User Interface....15
- Enterprise Management Window....15
  - Inheriting The System Settings....16
- Array Management Window....16
- Dell PowerVault Modular Disk Configuration Utility....17
- Other Information You May Need....17

2 About Your MD Series Storage Array....19
- Physical Disks, Virtual Disks, And Disk Groups....19
  - Physical Disks....19
  - Physical Disk States....19
  - Self-Monitoring Analysis And Reporting Technology....20
  - Virtual Disks And Disk Groups....20
  - Virtual Disk States....20
- Disk Pools....21
- Thin Virtual Disks....21
- RAID Levels....21
  - RAID Level Usage....22
  - RAID 0....22
  - RAID 1....22
  - RAID 5....22
  - RAID 6....22
  - RAID 10....23
- Segment Size....23
- Virtual Disk Operations....23
  - Virtual Disk Initialization....23
  - Consistency Check....23
  - Media Verification....24
  - Cycle Time....24
  - Virtual Disk Operations Limit....24
- Disk Group Operations....24
  - RAID Level Migration....24
  - Segment Size Migration....25

Virtual Disk Capacity Expansion....25
Disk Group Expansion....25
Disk Group Defragmentation....25
Disk Group Operations Limit....25
RAID Background Operations Priority....26
Virtual Disk Migration And Disk Roaming....26
Disk Migration....26
Disk Roaming....27
Host Server-To-Virtual Disk Mapping....27
Host Types....28
Advanced Features....28
Types Of Snapshot Functionality Supported....28
Snapshot Virtual Disks (Legacy)....29
Snapshot (Legacy) Repository Virtual Disk....29
Virtual Disk Copy....29
Virtual Disk Recovery....30
Using Snapshot And Virtual Disk Copy Together....30
Multi-Path Software....30
Preferred And Alternate Controllers And Paths....31
Virtual Disk Ownership....31
Load Balancing....31
Monitoring System Performance....32
Monitoring Performance Using The Tabular View....32
Monitoring Performance Using The Graphical View....33

3 Discovering And Managing Your Storage Array....35
Out-Of-Band Management....35
In-Band Management....35
Access Virtual Disk....36
Storage Arrays....36
Automatic Discovery Of Storage Arrays....36
Manual Addition Of A Storage Array....36
Setting Up Your Storage Array....37
Locating Storage Arrays....37
Naming Or Renaming Storage Arrays....38
Setting A Password....38
Adding Or Editing A Comment To An Existing Storage Array....39
Removing Storage Arrays....39
Enabling Premium Features....40
Displaying Failover Alert....40
Changing The Cache Settings On The Storage Array....40
Changing Expansion Enclosure ID Numbers....41

Changing The Enclosure Order....41
Configuring Alert Notifications....41
Configuring E-mail Alerts....42
Configuring SNMP Alerts....43
Battery Settings....44
Changing The Battery Settings....44
Setting The Storage Array RAID Controller Module Clocks....44

4 Using iSCSI....47
Changing The iSCSI Target Authentication....47
Entering Mutual Authentication Permissions....47
Creating CHAP Secrets....48
Initiator CHAP Secret....48
Target CHAP Secret....48
Valid Characters For CHAP Secrets....48
Changing The iSCSI Target Identification....49
Changing The iSCSI Target Discovery Settings....49
Configuring The iSCSI Host Ports....49
Advanced iSCSI Host Port Settings....50
Viewing Or Ending An iSCSI Session....51
Viewing iSCSI Statistics And Setting Baseline Statistics....51
Edit, Remove, Or Rename Host Topology....52

5 Event Monitor....53
Enabling Or Disabling The Event Monitor....53
Windows....53
Linux....53

6 About Your Host....55
Configuring Host Access....55
Using The Host Mappings Tab....55
Defining A Host....56
Removing Host Access....57
Managing Host Groups....57
Creating A Host Group....57
Adding A Host To A Host Group....58
Removing A Host From A Host Group....58
Moving A Host To A Different Host Group....58
Removing A Host Group....58
Host Topology....59
Starting Or Stopping The Host Context Agent....59
I/O Data Path Protection....59

Managing Host Port Identifiers....60

7 Disk Groups, Standard Virtual Disks, And Thin Virtual Disks....63
Creating Disk Groups And Virtual Disks....63
Creating Disk Groups....63
Locating A Disk Group....65
Creating Standard Virtual Disks....65
Changing The Virtual Disk Modification Priority....66
Changing The Virtual Disk Cache Settings....67
Changing The Segment Size Of A Virtual Disk....68
Changing The IO Type....68
Thin Virtual Disks....69
Advantages Of Thin Virtual Disks....69
Physical Vs Virtual Capacity On A Thin Virtual Disk....69
Thin Virtual Disk Requirements And Limitations....70
Thin Volume Attributes....70
Thin Virtual Disk States....71
Comparison—Types Of Virtual Disks And Copy Services....71
Rollback On Thin Virtual Disks....72
Initializing A Thin Virtual Disk....72
Changing A Thin Virtual Disk To A Standard Virtual Disk....75
Choosing An Appropriate Physical Disk Type....75
Physical Disk Security With Self Encrypting Disk....75
Creating A Security Key....77
Changing A Security Key....78
Saving A Security Key....79
Validate Security Key....80
Unlocking Secure Physical Disks....80
Erasing Secure Physical Disks....80
Configuring Hot Spare Physical Disks....80
Hot Spares And Rebuild....81
Global Hot Spares....82
Hot Spare Operation....82
Hot Spare Drive Protection....82
Enclosure Loss Protection....83
Drawer Loss Protection....83
Host-To-Virtual Disk Mapping....84
Creating Host-To-Virtual Disk Mappings....85
Modifying And Removing Host-To-Virtual Disk Mapping....86
Changing Controller Ownership Of The Virtual Disk....86
Removing Host-To-Virtual Disk Mapping....87
Changing The RAID Controller Module Ownership Of A Disk Group....87

Changing The RAID Level Of A Disk Group....88
Removing A Host-To-Virtual Disk Mapping Using Linux DMMP....88
Restricted Mappings....89
Storage Partitioning....90
Disk Group And Virtual Disk Expansion....91
Disk Group Expansion....91
Virtual Disk Expansion....91
Using Free Capacity....91
Using Unconfigured Capacity....92
Disk Group Migration....92
Export Disk Group....92
Import Disk Group....93
Storage Array Media Scan....94
Changing Media Scan Settings....94
Suspending The Media Scan....95

8 Disk Pools And Disk Pool Virtual Disks....97
Difference Between Disk Groups And Disk Pools....97
Disk Pool Restrictions....97
Creating A Disk Pool Manually....98
Automatically Managing The Unconfigured Capacity In Disk Pools....99
Locating Physical Disks In A Disk Pool....99
Renaming A Disk Pool....100
Configuring Alert Notifications For A Disk Pool....100
Adding Unassigned Physical Disks To A Disk Pool....101
Configuring The Preservation Capacity Of A Disk Pool....101
Changing The Modification Priority Of A Disk Pool....102
Changing The RAID Controller Module Ownership Of A Disk Pool....102
Checking Data Consistency....103
Deleting A Disk Pool....104
Viewing Storage Array Logical Components And Associated Physical Components....104
Secure Disk Pools....105
Changing Capacity On Existing Thin Virtual Disks....106
Creating A Thin Virtual Disk From A Disk Pool....106

9 Using SSD Cache....109
How SSD Cache Works....109
Benefits Of SSD Cache....109
Choosing SSD Cache Parameters....109
SSD Cache Restrictions....110
Creating An SSD Cache....110
Viewing Physical Components Associated With An SSD Cache....110

Locating Physical Disks In An SSD Cache....111
Adding Physical Disks To An SSD Cache....111
Removing Physical Disks From An SSD Cache....111
Suspending Or Resuming SSD Caching....112
Changing I/O Type In An SSD Cache....112
Renaming An SSD Cache....112
Deleting An SSD Cache....112
Using The Performance Modeling Tool....113

10 Premium Feature—Snapshot Virtual Disk....115
Snapshot Virtual Disk Vs. Snapshot Virtual Disk (Legacy)....115
Snapshot Images And Groups....115
Snapshot Virtual Disk Read/Write Properties....116
Snapshot Groups And Consistency Groups....116
Snapshot Groups....116
Snapshot Consistency Groups....117
Understanding Snapshot Repositories....117
Consistency Group Repositories....117
Ranking Repository Candidates....118
Using Snapshot Consistency Groups With Remote Replication....118
Creating Snapshot Images....118
Creating A Snapshot Image....119
Canceling A Pending Snapshot Image....119
Deleting A Snapshot Image....120
Scheduling Snapshot Images....120
Creating A Snapshot Schedule....121
Editing A Snapshot Schedule....121
Performing Snapshot Rollbacks....122
Snapshot Rollback Limitations....122
Starting A Snapshot Rollback....123
Resuming A Snapshot Image Rollback....123
Canceling A Snapshot Image Rollback....124
Viewing The Progress Of A Snapshot Rollback....124
Changing Snapshot Rollback Priority....125
Creating A Snapshot Group....125
Creating A Consistency Group Repository (Manually)....126
Changing Snapshot Group Settings....128
Renaming A Snapshot Group....128
Deleting A Snapshot Group....129
Converting Snapshots (Legacy) To A Snapshot Group....129
Conversion Limitations....129
Converting Snapshots (Legacy)....129

Converting A Snapshot Virtual Disk To Read-Write....130
Viewing Associated Physical Components Of An Individual Repository Virtual Disk....130
Creating A Consistency Group....131
Creating A Consistency Group Repository (Manually)....132
Renaming A Consistency Group....133
Deleting A Consistency Group....133
Changing The Settings Of A Consistency Group....134
Adding A Member Virtual Disk To A Consistency Group....134
Removing A Member Virtual Disk From A Consistency Group....135
Creating A Snapshot Virtual Disk Of A Snapshot Image....136
Snapshot Virtual Disk Limitations....136
Creating A Snapshot Virtual Disk....136
Creating A Snapshot Virtual Disk Repository....138
Changing The Settings Of A Snapshot Virtual Disk....139
Disabling A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....139
Re-creating A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....140
Renaming A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....141
Deleting A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....141
Creating A Consistency Group Snapshot Virtual Disk....142
Creating A Consistency Group Snapshot Virtual Disk Repository (Manually)....144
Disabling A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....145
Re-creating A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk....146
Changing The Modification Priority Of An Overall Repository Virtual Disk....146
Changing The Media Scan Setting Of An Overall Repository Virtual Disk....147
Changing The Pre-read Consistency Check Setting Of An Overall Repository Virtual Disk....147
Increasing The Capacity Of An Overall Repository....148
Decreasing The Capacity Of The Overall Repository....150
Performing A Revive Operation....151

11 Premium Feature—Snapshot Virtual Disks (Legacy)....153
Scheduling A Snapshot Virtual Disk....154
Common Reasons For Scheduling A Snapshot Virtual Disk....154
Guidelines for Creating Snapshot Schedules....154
Creating A Snapshot Virtual Disk Using The Simple Path....154
About The Simple Path....155
Preparing Host Servers To Create The Snapshot Using The Simple Path....155
Creating A Snapshot Virtual Disk Using The Advanced Path....156
About The Advanced Path....156
Preparing Host Servers To Create The Snapshot Using The Advanced Path....157
Creating The Snapshot Using The Advanced Path....158
Specifying Snapshot Virtual Disk Names....159
Snapshot Repository Capacity....160

Re-Creating Snapshot Virtual Disks....161
Disabling A Snapshot Virtual Disk....162
Preparing Host Servers To Re-Create A Snapshot Virtual Disk....162
Re-Creating A Snapshot Virtual Disk....163

12 Premium Feature—Virtual Disk Copy....165
Using Virtual Disk Copy With Snapshot Or Snapshot (Legacy) Premium Feature....165
Types Of Virtual Disk Copies....166
Offline Copy....166
Online Copy....166
Creating A Virtual Disk Copy For An MSCS Shared Disk....166
Virtual Disk Read/Write Permissions....167
Virtual Disk Copy Restrictions....167
Creating A Virtual Disk Copy....168
Setting Read/Write Permissions On Target Virtual Disk....168
Before You Begin....168
Virtual Disk Copy And Modification Operations....168
Create Copy Wizard....169
Failed Virtual Disk Copy....169
Preferred RAID Controller Module Ownership....169
Failed RAID Controller Module....169
Copy Manager....170
Copying The Virtual Disk....170
Storage Array Performance During Virtual Disk Copy....171
Setting Copy Priority....171
Stopping A Virtual Disk Copy....171
Recopying A Virtual Disk....172
Preparing Host Servers To Recopy A Virtual Disk....172
Recopying The Virtual Disk....173
Removing Copy Pairs....173

13 Device Mapper Multipath For Linux....175
Overview....175
Using DM Multipathing Devices....175
Prerequisites....175
Device Mapper Configuration Steps....176
Scan For Newly Added Virtual Disks....176
Display The Multipath Device Topology Using The Multipath Command....177
Create A New fdisk Partition On A Multipath Device Node....177
Add A New Partition To Device Mapper....178
Create A File System On A Device Mapper Partition....178
Mount A Device Mapper Partition....178

Ready For Use....178
Linux Host Server Reboot Best Practices....178
Important Information About Special Partitions....179
Limitations And Known Issues....179
Troubleshooting....180

14 Configuring Asymmetric Logical Unit Access....181
ALUA Performance Considerations....181
Automatic Transfer Of Ownership....181
Native ALUA Support On Microsoft Windows And Linux....181
Enabling ALUA On VMware ESX/ESXi....181
Manually Adding SATP Rule In ESX 4.1 (Update 1 And Update 2)....181
Manually Adding SATP Rule In ESX/ESXi 5.0....182
Verifying ALUA On VMware ESX/ESXi....182
Verifying If Host Server Is Using ALUA For MD Storage Array....182
Setting Round-Robin Load Balancing Policy On ESX/ESXi-Based Storage Arrays....183

15 Premium Feature—Remote Replication....185
About Asynchronous Remote Replication....185
Remote Replicated Pairs And Replication Repositories....185
Types Of Remote Replication....186
Differences Between Remote Replication Features....186
Upgrading To Asynchronous Remote Replication From Remote Replication (Legacy)....186
Remote Replication Requirements And Restrictions....187
Restrictions On Using Remote Replication....187
Setting Up Remote Replication....187
Activating Remote Replication Premium Features....188
Deactivating Remote Replication....188
Remote Replication Groups....189
Purpose Of A Remote Replication Group....189
Remote Replication Group Requirements And Guidelines....189
Creating A Remote Replication Group....189
Replicated Pairs....190
Guidelines for Choosing Virtual Disks in a Replicated Pair....190
Guidelines For Choosing Virtual Disks In A Replicated Pair....190
Creating Replicated Pairs....191
Removing A Replicated Pair From A Remote Replication Group....192

16 Premium Feature—Remote Replication (Legacy)....193
Overview....193
Switchable Host Access Configuration Required With Remote Replication (Legacy)....193
Activating Remote Replication (Legacy)....194

Replication Repository Virtual Disks....194
RAID Levels For Replication Repository Virtual Disks....194
Primary And Secondary Virtual Disk Pairs....194
Using Remote Replication (Legacy) With Other Features....195
Storage Partitioning With Remote Replication (Legacy)....195
Snapshot Virtual Disk With Remote Replication (Legacy)....195
Virtual Disk Copy With Remote Replication (Legacy)....195
Virtual Disk Expansion With Remote Replication (Legacy)....196
Required Switch Zoning Configurations....196
Zoning Guidelines for Remote Replication (Legacy)....196
Switch Cascading....197
Journaling File Systems And Remote Replication (Legacy)....197
Prerequisites For Setting Up And Managing Remote Replication (Legacy)....197
Activating The Remote Replication (Legacy) Premium Feature And Creating Replication Virtual Disks....197
Activating And Creating Replication Repository Virtual Disks From The Storage Array....197
Activating And Creating Replication Repository Virtual Disks From An Existing Disk Group....198
Creating A Remote Replication (Legacy)....198
Selecting The Secondary Virtual Disk....199
Setting The Write Mode....199
Setting Synchronization Priority And Synchronization Method....200
Completing The Remote Replication (Legacy)....200
RAID Controller Module Ownership/Preferred Path....201
Viewing Information About A Remote Replication (Legacy) Or Replication Repository (Legacy) Virtual Disk....201
Viewing The Storage Array Profile....201
Viewing The Properties Pane....202
Viewing Logical Elements Of The Secondary Virtual Disk....202
Viewing The Physical Components Or Logical Elements Of The Primary Virtual Disk....202
Virtual Disk Status Icons....203
Changing Write Mode And Consistency Group Membership....203
Resynchronizing Virtual Disks....204
Normally Synchronized Virtual Disks....204
Changing Synchronization Priority And Method....205
Unsynchronized Virtual Disks....205
Automatically Resynchronizing Virtual Disks....206
Manually Resynchronizing Virtual Disks....206
Reversing Roles Between The Primary And Secondary Virtual Disks....206
Promoting The Secondary Virtual Disk Or Demoting The Primary Virtual Disk....207
Suspending A Remote Replication (Legacy)....207
Resuming A Remote Replication (Legacy)....207
About Resumed Remote Replications....208
Testing Communication Between The Primary And Secondary Virtual Disks....208
Deleting A Virtual Disk From A Replicated Pair In A Storage Array....209

17 Management Firmware Downloads....213
Downloading RAID Controller And NVSRAM Packages....213
Downloading Both RAID Controller And NVSRAM Firmware....213
Downloading Only NVSRAM Firmware....215
Downloading Physical Disk Firmware....216
Downloading MD3060e Series Expansion Module EMM Firmware....217
Self-Monitoring Analysis And Reporting Technology (SMART)....218
Media Errors And Unreadable Sectors....218
18 Firmware Inventory....219
Viewing The Firmware Inventory....219
19 System Interfaces....221
Virtual Disk Service....221
Volume Shadow-Copy Service....221
20 Storage Array Software....223
Start-Up Routine....223
Device Health Conditions....223
Trace Buffers....225
Retrieving Trace Buffers....226
Collecting Physical Disk Data....226
Creating A Support Data Collection Schedule....227
Suspending Or Resuming A Support Data Collection Schedule....227
Removing A Support Data Collection Schedule....228
Event Log....228
Viewing The Event Log....228
Recovery Guru....229
Storage Array Profile....229
Viewing The Physical Associations....230
Recovering From An Unresponsive Storage Array Condition....230
Locating A Physical Disk....232
Locating An Expansion Enclosure....232
Capturing The State Information....233
SMrepassist Utility....233
Unidentified Devices....234
Recovering From An Unidentified Storage Array....234
Starting Or Restarting The Host Context Agent Software....235
Starting The SMagent Software In Windows....235
Starting The SMagent Software In Linux....235

21 Getting Help.....237
Contacting Dell.....237

1

Introduction

CAUTION: See the Safety, Environmental, and Regulatory Information document for important safety information before following any procedures listed in this document.

This document familiarizes you with the functions of the Dell PowerVault MD Series storage array. The document is organized according to the tasks that you must complete after receiving your storage array.

Dell PowerVault Modular Disk Storage Manager

Dell PowerVault Modular Disk Storage Manager (MD Storage Manager) is a graphical user interface (GUI) application used to configure and manage one or more MD Series storage arrays. The MD Storage Manager software is located on the MD Series resource DVD.

User Interface

The Storage Manager screen is divided into two primary windows:

- Enterprise Management Window (EMW) — The EMW provides high-level management of multiple storage arrays. You can launch the Array Management Windows for the storage arrays from the EMW.
- Array Management Window (AMW) — The AMW provides management functions for a single storage array.

The EMW and the AMW consist of the following:

- The title bar at the top of the window — Shows the name of the application.
- The menu bar, beneath the title bar — You can select menu options from the menu bar to perform tasks on a storage array.
- The toolbar, beneath the menu bar — You can select options in the toolbar to perform tasks on a storage array.

  NOTE: The toolbar is available only in the EMW.
- The tabs, beneath the toolbar — Tabs are used to group the tasks that you can perform on a storage array.
- The status bar, beneath the tabs — The status bar shows status messages and status icons related to the storage array.

  NOTE: By default, the toolbar and status bar are not displayed. To view the toolbar or the status bar, select **View** → **Toolbar** or **View** → **Status Bar**, respectively.

Enterprise Management Window

The EMW provides high-level management of storage arrays. When you start the MD Storage Manager, the EMW is displayed. The EMW has the:

- **Devices** tab — Provides information about discovered storage arrays.
- **Setup** tab — Presents the initial setup tasks that guide you through adding storage arrays and configuring alerts.

The **Devices** tab has a Tree view on the left side of the window that shows discovered storage arrays, unidentified storage arrays, and the status conditions for the storage arrays. Discovered storage arrays are managed by the MD

Storage Manager. Unidentified storage arrays are available to the MD Storage Manager but not configured for management. The right side of the **Devices** tab has a Table view that shows detailed information for the selected storage array.

In the EMW, you can:

- Discover hosts and managed storage arrays on the local sub-network.
- Manually add and remove hosts and storage arrays.
- Blink or locate the storage arrays.
- Name or rename discovered storage arrays.
- Add comments for a storage array in the Table view.
- Store your EMW view preferences and configuration data in local configuration files. The next time you open the EMW, data from the local configuration files is used to show customized view and preferences.
- Monitor the status of managed storage arrays and indicate status using appropriate icons.
- Add or remove management connections.
- Configure alert notifications for all selected storage arrays through e-mail or SNMP traps.
- Report critical events to the configured alert destinations.
- Launch the AMW for a selected storage array.
- Run a script to perform batch management tasks on specific storage arrays.
- Import the operating system theme settings into the MD Storage Manager.
- Upgrade firmware on multiple storage arrays concurrently.
- Obtain information about the firmware inventory including the version of the RAID controller modules, physical disks, and the enclosure management modules (EMMs) in the storage array.

Inheriting The System Settings

Use the **Inherit System Settings** option to import the operating system theme settings into the MD Storage Manager. Importing system theme settings affects the font type, font size, color, and contrast in the MD Storage Manager.

1. From the EMW, open the **Inherit System Settings** window in one of these ways:
   - Select **Tools → Inherit System Settings.**
   - Select the **Setup** tab, and under **Accessibility**, click **Inherit System Settings**.
2. Select **Inherit system settings for color and font**.
3. Click **OK**.

Array Management Window

You can launch the AMW from the EMW. The AMW provides management functions for a single storage array. You can have multiple AMWs open simultaneously to manage different storage arrays.

To launch the AMW:

1. In the EMW, on the **Devices** tab, right-click on the relevant storage array.
   The context menu for the selected storage is displayed.
2. In the context menu, select **Manage Storage Array**.
   The AMW for the selected storage array is displayed.

NOTE: You can also launch the AMW by:

- Double-clicking on a storage array displayed in the **Devices** tab of the EMW.
- Selecting a storage array displayed in the **Devices** tab of the EMW, and then selecting **Tools → Manage Storage Array.**

The AMW has the following tabs:

- **Summary** tab — You can view the following information about the storage array:
  - Status
  - Hardware
  - Storage and copy services
  - Hosts and mappings
  - Information on storage capacity
  - Premium features
- **Storage & Copy Services** tab — You can view and manage the organization of the storage array by virtual disks, disk groups, free capacity nodes, and any unconfigured capacity for the storage array.
- **Host Mappings** tab — You can define the hosts, host groups, and host ports. You can change the mappings to grant virtual disk access to host groups and hosts and create storage partitions.
- **Hardware** tab — You can view and manage the physical components of the storage array.
- **Setup** tab — Shows a list of initial setup tasks for the storage array.

In the AMW, you can:

- Select storage array options — For example, renaming a storage array, changing a password, or enabling a background media scan.
- Configure virtual disks and disk pools from the storage array capacity, define hosts and host groups, and grant host or host group access to sets of virtual disks called storage partitions.
- Monitor the health of storage array components and report detailed status using applicable icons.
- Perform recovery procedures for a failed logical component or a failed hardware component.
- View the Event Log for a storage array.
- View profile information about hardware components, such as RAID controller modules and physical disks.
- Manage RAID controller modules — For example, changing ownership of virtual disks or placing a RAID controller module online or offline.
- Manage physical disks — For example, assignment of hot spares and locating the physical disk.
- Monitor storage array performance.

Dell PowerVault Modular Disk Configuration Utility

NOTE: Dell PowerVault Modular Disk Configuration Utility (MDCU) is supported only on MD Series storage arrays that use the iSCSI protocol.

MDCU is an iSCSI Configuration Wizard that can be used in conjunction with MD Storage Manager to simplify the configuration of iSCSI connections. The MDCU software is available on the MD Series resource media.

Other Information You May Need

WARNING: See the safety and regulatory information that shipped with your system. Warranty information may be included within this document or as a separate document.

NOTE: All the documents, unless specified otherwise, are available at **dell.com/support/manuals**.

- The *Getting Started Guide* provides an overview of setting up and cabling your storage array.
- The *Deployment Guide* provides installation and configuration instructions for both software and hardware.
- The *Owner's Manual* provides information about system features and describes how to troubleshoot the system and install or replace system components.
- The *CLI Guide* provides information about using the command line interface (CLI).
- The MD Series resource media contains all system management tools.
- The *Dell PowerVault MD3260, MD3260i, MD3660i, MD3660f and MD3060e Series Storage Arrays Support Matrix* provides information on supported software and hardware for MD systems.
- Information Updates or readme files are included to provide last-minute updates to the enclosure or documentation or advanced technical reference material intended for experienced users or technicians.
- For video resources on PowerVault MD storage arrays, go to **dell.com/techcenter**.
- For the full name of an abbreviation or acronym used in this document, see the Glossary at **dell.com/support/manuals**.

NOTE: Always check for updates on **dell.com/support/manuals** and read the updates first because they often supersede information in other documents.

2

About Your MD Series Storage Array

This chapter describes the storage array concepts, which help in configuring and operating the Dell MD Series storage arrays.

Physical Disks, Virtual Disks, And Disk Groups

Physical disks in your storage array provide the physical storage capacity for your data. Before you can begin writing data to the storage array, you must configure the physical storage capacity into logical components, called disk groups and virtual disks.

A disk group is a set of physical disks upon which multiple virtual disks are created. The maximum number of physical disks supported in a disk group is 96 disks for RAID 0, RAID 1, and RAID 10, and 30 drives for RAID 5 and RAID 6. You can create disk groups from unconfigured capacity on your storage array.

A virtual disk is a partition in a disk group that is made up of contiguous data segments of the physical disks in the disk group. A virtual disk consists of data segments from all physical disks in the disk group.

All virtual disks in a disk group support the same RAID level. The storage array supports up to 255 virtual disks (minimum size of 10 MB each) that can be assigned to host servers. Each virtual disk is assigned a Logical Unit Number (LUN) that is recognized by the host operating system.

Virtual disks and disk groups are set up according to how you plan to organize your data. For example, you can have one virtual disk for inventory, a second virtual disk for financial and tax information, and so on.

Physical Disks

Only Dell supported physical disks are supported in the storage array. If the storage array detects unsupported physical disks, it marks the disk as unsupported and the physical disk becomes unavailable for all operations.

For the list of supported physical disks, see the Support Matrix at **dell.com/support/manuals**.

Physical Disk States

The following describes the various states of the physical disk, which are recognized by the storage array and reported in the MD Storage Manager.

Status	Mode	Description
Optimal	Assigned	The physical disk in the indicated slot is configured as part of a disk group.
Optimal	Unassigned	The physical disk in the indicated slot is unused and available to be configured.
Optimal	Hot Spare Standby	The physical disk in the indicated slot is configured as a hot spare.

Status	Mode	Description
Optimal	Hot Spare in use	The physical disk in the indicated slot is in use as a hot spare within a disk group.
Failed	Assigned, Unassigned, Hot Spare in use, or Hot Spare Standby	The physical disk in the indicated slot has failed because of an unrecoverable error, an incorrect drive type or drive size, or by its operational state being set to failed.
Replaced	Assigned	The physical disk in the indicated slot has been replaced and is ready to be, or is actively being, configured into a disk group.
Pending Failure	Assigned, Unassigned, Hot Spare in use, or Hot Spare Standby	A Self-Monitoring Analysis and Reporting Technology (SMART) error has been detected on the physical disk in the indicated slot.
Offline	Not applicable	The physical disk has either been spun down or had a rebuild aborted by user request.
Identify	Assigned, Unassigned, Hot Spare in use, or Hot Spare Standby	The physical disk is being identified.

Self-Monitoring Analysis And Reporting Technology

Self-Monitoring Analysis and Reporting Technology (SMART) monitors the internal performance of all physical disk components to detect faults indicating the potential for physical disk failure. SMART uses this information to report whether failure is imminent so that a physical disk can be replaced before failure occurs. The storage array monitors all attached drives and notifies you when a predicted failure is reported by a physical disk.

Virtual Disks And Disk Groups

When configuring a storage array, you must:

- Organize the physical disks into disk groups.
- Create virtual disks within these disk groups.
- Provide host server access.
- Create mappings to associate the virtual disks with the host servers.

  NOTE: Host server access must be created before mapping virtual disks.

Disk groups are always created in the unconfigured capacity of a storage array. Unconfigured capacity is the available physical disk space not already assigned in the storage array.

Virtual disks are created within the free capacity of a disk group. Free capacity is the space in a disk group that has not been assigned to a virtual disk.

Virtual Disk States

The following table describes the various states of the virtual disk, recognized by the storage array.

Table 1. RAID Controller Virtual Disk States

State	Description
Optimal	The virtual disk contains physical disks that are online.
Degraded	The virtual disk with a redundant RAID level contains an inaccessible physical disk. The system can still function properly, but performance may be affected and additional disk failures may result in data loss.
Offline	A virtual disk with one or more member disks in an inaccessible (failed, missing, or offline) state. Data on the virtual disk is no longer accessible.
Force online	The storage array forces a virtual disk that is in an **Offline** state to an **Optimal** state. If all the member physical disks are not available, the storage array forces the virtual disk to a **Degraded** state. The storage array can force a virtual disk to an **Online** state only when a sufficient number of physical disks are available to support the virtual disk.

Disk Pools

Disk pooling allows you to distribute data from each virtual disk randomly across a set of physical disks. Although there is no limit on the maximum number of physical disks that can comprise a disk pool, each disk pool must have a minimum of 11 physical disks. Additionally, the disk pool cannot contain more physical disks than the maximum limit for each storage array.

Thin Virtual Disks

Thin virtual disks can be created from an existing disk pool. Creating thin virtual disks allows you to set up a large virtual space, but only use the actual physical space as you need it.

RAID Levels

RAID levels determine the way in which data is written to physical disks. Different RAID levels provide different levels of accessibility, redundancy, and capacity.

Using multiple physical disks has the following advantages over using a single physical disk:

- Placing data on multiple physical disks (striping) allows input/output (I/O) operations to occur simultaneously and improve performance.
- Storing redundant data on multiple physical disks using mirroring or parity supports reconstruction of lost data if an error occurs, even if that error is the failure of a physical disk.

Each RAID level provides different performance and protection. You must select a RAID level based on the type of application, access, fault tolerance, and data you are storing.

The storage array supports RAID levels 0, 1, 5, 6, and 10. The maximum and minimum number of physical disks that can be used in a disk group depends on the RAID level:

- 192 for RAID levels 0, 1, and 10
- 30 for RAID 5 and 6

RAID Level Usage

To ensure best performance, you must select an optimal RAID level when you create a system physical disk. The optimal RAID level for your disk array depends on:

- Number of physical disks in the disk array
- Capacity of the physical disks in the disk array
- Need for redundant access to the data (fault tolerance)
- Disk performance requirements

RAID 0

CAUTION: Do not attempt to create virtual disk groups exceeding 120 physical disks in a RAID 0 configuration even if premium feature is activated on your storage array. Exceeding the 120-physical disk limit may cause your storage array to be unstable.

RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy. RAID 0 breaks the data down into segments and writes each segment to a separate physical disk. I/O performance is greatly improved by spreading the I/O load across many physical disks. Although it offers the best performance of any RAID level, RAID 0 lacks data redundancy. Choose this option only for non-critical data, because failure of one physical disk results in the loss of all data. Examples of RAID 0 applications include video editing, image editing, prepress applications, or any application that requires high bandwidth.

RAID 1

RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. RAID 1 offers fast performance and the best data availability, but also the highest disk overhead. RAID 1 is recommended for small databases or other applications that do not require large capacity. For example, accounting, payroll, or financial applications. RAID 1 provides full data redundancy.

RAID 5

RAID 5 uses parity and striping data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access. RAID 5 is a versatile RAID level and is suited for multi-user environments where typical I/O size is small and there is a high proportion of read activity such as file, application, database, web, e-mail, news, and intranet servers.

RAID 6

RAID 6 is similar to RAID 5 but provides an additional parity disk for better redundancy. RAID 6 is the most versatile RAID level and is suited for multi-user environments where typical I/O size is small and there is a high proportion of read activity. RAID 6 is recommended when large size physical disks are used or large number of physical disks are used in a disk group.

RAID 10

CAUTION: Do not attempt to create virtual disk groups exceeding 120 physical disks in a RAID 10 configuration even if premium feature is activated on your storage array. Exceeding the 120-physical disk limit may cause your storage array to be unstable.

RAID 10, a combination of RAID 1 and RAID 0, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy. Utilizing an even number of physical disks (four or more) creates a RAID level 10 disk group and/or virtual disk. Because RAID levels 1 and 10 use disk mirroring, half of the capacity of the physical disks is utilized for mirroring. This leaves the remaining half of the physical disk capacity for actual storage. RAID 10 is automatically used when a RAID level of 1 is chosen with four or more physical disks. RAID 10 works well for medium-sized databases or any environment that requires high performance and fault tolerance and moderate-to-medium capacity.

Segment Size

Disk striping enables data to be written across multiple physical disks. Disk striping enhances performance because striped disks are accessed simultaneously.

The segment size or stripe element size specifies the size of data in a stripe written to a single disk. The storage array supports stripe element sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, and 256 KB. The default stripe element size is 128 KB.

Stripe width, or depth, refers to the number of disks involved in an array where striping is implemented. For example, a four-disk group with disk striping has a stripe width of four.

NOTE: Although disk striping delivers excellent performance, striping alone does not provide data redundancy.

Virtual Disk Operations

Virtual Disk Initialization

Every virtual disk must be initialized. Initialization can be done in the foreground or the background. A maximum of four virtual disks can be initialized concurrently on each RAID controller module.

- Background initialization — The storage array executes a background initialization when the virtual disk is created to establish parity, while allowing full host server access to the virtual disks. Background initialization does not run on RAID 0 virtual disks. The background initialization rate is controlled by MD Storage Manager. To change the rate of background initialization, you must stop any existing background initialization. The rate change is implemented when the background initialization restarts automatically.
- Foreground Initialization — The storage array executes a background initialization when the virtual disk is created to establish parity, while allowing full host server access to the virtual disks. Background initialization does not run on RAID 0 virtual disks. The background initialization rate is controlled by MD Storage Manager. To change the rate of background initialization, you must stop any existing background initialization. The rate change is implemented when the background initialization restarts automatically.

Consistency Check

A consistency check verifies the correctness of data in a redundant array (RAID levels 1, 5, 6, and 10). For example, in a system with parity, checking consistency involves computing the data on one physical disk and comparing the results to the contents of the parity physical disk.

A consistency check is similar to a background initialization. The difference is that background initialization cannot be started or stopped manually, while consistency check can.

NOTE: It is recommended that you run data consistency checks on a redundant array at least once a month. This allows detection and automatic replacement of unreadable sectors. Finding an unreadable sector during a rebuild of a failed physical disk is a serious problem, because the system does not have the redundancy to recover the data.

Media Verification

Another background task performed by the storage array is media verification of all configured physical disks in a disk group. The storage array uses the Read operation to perform verification on the space configured in virtual disks and the space reserved for the metadata.

Cycle Time

The media verification operation runs only on selected disk groups, independent of other disk groups. Cycle time is the time taken to complete verification of the metadata region of the disk group and all virtual disks in the disk group for which media verification is configured. The next cycle for a disk group starts automatically when the current cycle completes. You can set the cycle time for a media verification operation between 1 and 30 days. The storage controller throttles the media verification I/O accesses to disks based on the cycle time.

The storage array tracks the cycle for each disk group independent of other disk groups on the controller and creates a checkpoint. If the media verification operation on a disk group is preempted or blocked by another operation on the disk group, the storage array resumes after the current cycle. If the media verification process on a disk group is stopped due to a RAID controller module restart, the storage array resumes the process from the last checkpoint.

Virtual Disk Operations Limit

The maximum number of active, concurrent virtual disk processes per RAID controller module installed in the storage array is four. This limit is applied to the following virtual disk processes:

- Background initialization
- Foreground initialization
- Consistency check
- Rebuild
- Copy back

If a redundant RAID controller module fails with existing virtual disk processes, the processes on the failed controller are transferred to the peer controller. A transferred process is placed in a suspended state if there are four active processes on the peer controller. The suspended processes are resumed on the peer controller when the number of active processes falls below four.

Disk Group Operations

RAID Level Migration

You can migrate from one RAID level to another depending on your requirements. For example, fault-tolerant characteristics can be added to a stripe set (RAID 0) by converting it to a RAID 5 set. The MD Storage Manager provides information about RAID attributes to assist you in selecting the appropriate RAID level. You can perform a RAID level migration while the system is still running and without rebooting, which maintains data availability.

Segment Size Migration

Segment size refers to the amount of data (in kilobytes) that the storage array writes on a physical disk in a virtual disk before writing data on the next physical disk. Valid values for the segment size are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, and 256 KB.

Dynamic segment size migration enables the segment size of a given virtual disk to be changed. A default segment size is set when the virtual disk is created, based on such factors as the RAID level and expected usage. You can change the default value if segment size usage does not match your needs.

When considering a segment size change, two scenarios illustrate different approaches to the limitations:

- If I/O activity stretches beyond the segment size, you can increase it to reduce the number of disks required for a single I/O. Using a single physical disk for a single request frees disks to service other requests, especially when you have multiple users accessing a database or storage environment.
- If you use the virtual disk in a single-user, large I/O environment (such as for multimedia application storage), performance can be optimized when a single I/O request is serviced with a single data stripe (the segment size multiplied by the number of physical disks in the disk group used for data storage). In this case, multiple disks are used for the same request, but each disk is only accessed once.

Virtual Disk Capacity Expansion

When you configure a virtual disk, you select a capacity based on the amount of data you expect to store. However, you may need to increase the virtual disk capacity for a standard virtual disk by adding free capacity to the disk group. This creates more unused space for new virtual disks or to expand existing virtual disks.

Disk Group Expansion

Because the storage array supports hot-swappable physical disks, you can add two physical disks at a time for each disk group while the storage array remains online. Data remains accessible on virtual disk groups, virtual disks, and physical disks throughout the operation. The data and increased unused free space are dynamically redistributed across the disk group. RAID characteristics are also reapplied to the disk group as a whole.

Disk Group Defragmentation

Defragmenting consolidates the free capacity in the disk group into one contiguous area. Defragmentation does not change the way in which the data is stored on the virtual disks.

Disk Group Operations Limit

The maximum number of active, concurrent disk group processes per installed RAID controller module is one. This limit is applied to the following disk group processes:

- Virtual disk RAID level migration
- Segment size migration
- Virtual disk capacity expansion
- Disk group expansion
- Disk group defragmentation

If a redundant RAID controller module fails with an existing disk group process, the process on the failed controller is transferred to the peer controller. A transferred process is placed in a suspended state if there is an active disk group

process on the peer controller. The suspended processes are resumed when the active process on the peer controller completes or is stopped.

NOTE: If you try to start a disk group process on a controller that does not have an existing active process, the start attempt fails if the first virtual disk in the disk group is owned by the other controller and there is an active process on the other controller.

RAID Background Operations Priority

The storage array supports a common configurable priority for the following RAID operations:

- Background initialization
- Rebuild
- Copy back
- Virtual disk capacity expansion
- Raid level migration
- Segment size migration
- Disk group expansion
- Disk group defragmentation

The priority of each of these operations can be changed to address performance requirements of the environment in which the operations are to be executed.

NOTE: Setting a high priority level impacts storage array performance. It is not advisable to set priority levels at the maximum level. Priority must also be assessed in terms of impact to host server access and time to complete an operation. For example, the longer a rebuild of a degraded virtual disk takes, the greater the risk for potential secondary disk failure.

Virtual Disk Migration And Disk Roaming

Virtual disk migration is moving a virtual disk or a hot spare from one array to another by detaching the physical disks and re-attaching them to the new array. Disk roaming is moving a physical disk from one slot to another on the same array.

Disk Migration

You can move virtual disks from one array to another without taking the target array offline. However, the disk group being migrated must be offline prior to performing the disk migration. If the disk group is not offline prior to migration, the source array holding the physical and virtual disks within the disk group marks them as missing. However, the disk groups themselves migrate to the target array.

An array can import a virtual disk only if it is in an optimal state. You can move virtual disks that are part of a disk group only if all members of the disk group are being migrated. The virtual disks automatically become available after the target array has finished importing all the disks in the disk group.

When you migrate a physical disk or a disk group from:

- One MD storage array to another MD storage array of the same type (for example, from an MD3260 storage array to another MD3260 storage array), the MD storage array you migrate to, recognizes any data structures and/or metadata you had in place on the migrating MD storage array.
- Any storage array different from the MD storage array you migrate to (for example, from an MD3260 storage array to an MD3260i storage array), the receiving storage array (MD3260i storage array in the example) does not

recognize the migrating metadata and that data is lost. In this case, the receiving storage array initializes the physical disks and marks them as unconfigured capacity.

NOTE: Only disk groups and associated virtual disks with all member physical disks present can be migrated from one storage array to another. It is recommended that you only migrate disk groups that have all their associated member virtual disks in an optimal state.

NOTE: The number of physical disks and virtual disks that a storage array supports limits the scope of the migration.

Use either of the following methods to move disk groups and virtual disks:

- Hot virtual disk migration — Disk migration with the destination storage array power turned on.
- Cold virtual disk migration — Disk migration with the destination storage array power turned off.

NOTE: To ensure that the migrating disk groups and virtual disks are correctly recognized when the target storage array has an existing physical disk, use hot virtual disk migration.

When attempting virtual disk migration, follow these recommendations:

- Moving physical disks to the destination array for migration — When inserting drives into the destination storage array during hot virtual disk migration, wait for the inserted physical disk to be displayed in the MD Storage Manager, or wait for 30 seconds (whichever occurs first), before inserting the next physical disk.

  WARNING: Without the interval between drive insertions, the storage array may become unstable and manageability may be temporarily lost.
- Migrating virtual disks from multiple storage arrays into a single storage array — When migrating virtual disks from multiple or different storage arrays into a single destination storage array, move all of the physical disks from the same storage array as a set into the new destination storage array. Ensure that all of the physical disks from a storage array are migrated to the destination storage array before starting migration from the next storage array.

  NOTE: If the drive modules are not moved as a set to the destination storage array, the newly relocated disk groups may not be accessible.
- Migrating virtual disks to a storage array with no existing physical disks — Turn off the destination storage array, when migrating disk groups or a complete set of physical disks from a storage array to another storage array that has no existing physical disks. After the destination storage array has been turned on and has successfully recognized the newly migrated physical disks, migration operations can continue.

  NOTE: Disk groups from multiple storage arrays must not be migrated at the same time to a storage array that has no existing physical disks. Use cold virtual disk migration for the disk groups from one storage array.
- Enabling premium features before migration — Before migrating disk groups and virtual disks, enable the required premium features on the destination storage array. If a disk group is migrated from a storage array that has a premium feature enabled and the destination array does not have this feature enabled, an **Out of Compliance** error message can be generated.

Disk Roaming

You can move physical disks within an array. The RAID controller module automatically recognizes the relocated physical disks and logically places them in the proper virtual disks that are part of the disk group. Disk roaming is permitted when the RAID controller module is either online or powered off.

NOTE: The disk group must be exported before moving the physical disks.

Host Server-To-Virtual Disk Mapping

The host server attached to a storage array accesses various virtual disks on the storage array through its host ports. Specific virtual disk-to-LUN mappings to an individual host server can be defined. In addition, the host server can be

part of a host group that shares access to one or more virtual disks. You can manually configure a host server-to-virtual disk mapping. When you configure host server-to-virtual disk mapping, consider these guidelines:

- You can define one host server-to-virtual disk mapping for each virtual disk in the storage array.
- Host server-to-virtual disk mappings are shared between RAID controller modules in the storage array.
- A unique LUN must be used by a host group or host server to access a virtual disk.
- Not every operating system has the same number of LUNs available for use.

Host Types

A host server is a server that accesses a storage array. Host servers are mapped to the virtual disks and use one or more iSCSI initiator ports. Host servers have the following attributes:

- Host name — A name that uniquely identifies the host server.
- Host group (used in Cluster solutions only) — Two or more host servers associated together to share access to the same virtual disks.

  NOTE: This host group is a logical entity you can create in the MD Storage Manager. All host servers in a host group must be running the same operating system.
- Host type — The operating system running on the host server.

Advanced Features

The RAID enclosure supports several advanced features:

- Virtual Disk Snapshots.
- Virtual Disk Copy.
- Remote replication (standard and legacy-based). For more information, see Premium Feature - Remote Replication and Premium Feature- Remote Replication (Legacy).

NOTE: The premium features listed above must be activated separately. If you have purchased these features, an activation card is supplied that contains instructions for enabling this functionality.

Types Of Snapshot Functionality Supported

The following types of virtual disk snapshot premium features are supported on the MD storage array:

- Snapshot Virtual Disks using multiple point-in-time (PiT) groups — This feature also supports snapshot groups, snapshot images, and consistency groups.
- Snapshot Virtual Disks (Legacy) using a separate repository for each snapshot

For more information, see Premium Feature---Snapshot Virtual Disk and Premium Feature—Snapshot Virtual Disks (Legacy).

Snapshot Virtual Disks, Snapshot Images, And Snapshot Groups

A snapshot image is a logical image of the content of an associated base virtual disk created at a specific point-in-time. This type of image is not directly readable or writable to a host since the snapshot image is used to save data from the base virtual disk only. To allow the host to access a copy of the data in a snapshot image, you must create a snapshot virtual disk. This snapshot virtual disk contains its own repository, which is used to save subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image.

Snapshot images can be created manually or automatically by establishing a schedule that defines the date and time you want to create the snapshot image. The following objects can be included in a snapshot image:

- Standard virtual disks
- Thin provisioned virtual disks
- Consistency groups

To create a snapshot image, you must first create a snapshot group and reserve snapshot repository space for the virtual disk. The repository space is based on a percentage of the current virtual disk reserve.

You can delete the oldest snapshot image in a snapshot group either manually or you can automate the process by enabling the **Auto-Delete** setting for the snapshot group. When a snapshot image is deleted, its definition is removed from the system, and the space occupied by the snapshot image in the repository is released and made available for reuse within the snapshot group.

Snapshot Virtual Disks (Legacy)

A snapshot is a point-in-time image of a virtual disk. The snapshot provides an image of the virtual disk at the time the snapshot was created. You create a snapshot so that an application (for example, a backup application) can access the snapshot and read the data while the source virtual disk remains online and user-accessible. When the backup is completed, the snapshot virtual disk is no longer needed. You can create up to four snapshots per virtual disk.

Snapshots are used to recover previous versions of files that have changed since the snapshot was taken. Snapshots are implemented using a copy on write algorithm, which makes a backup copy of data the instant a write occurs to the virtual disk. Data on a virtual disk is copied to the snapshot repository before it is modified. Snapshots are instantaneous and take up less overhead than a full physical copy process.

Snapshot (Legacy) Repository Virtual Disk

When you create a snapshot virtual disk, it automatically creates a snapshot repository virtual disk. A snapshot repository is a virtual disk created in the storage array as a resource for a snapshot virtual disk. A snapshot repository virtual disk contains snapshot virtual disk metadata and copy-on-write data for a particular snapshot virtual disk. The repository supports one snapshot only.

You cannot select a snapshot repository virtual disk as a source virtual disk or as a target virtual disk in a virtual disk copy. If you select a Snapshot source virtual disk as the target virtual disk of a virtual disk copy, you must disable all snapshot virtual disks associated with the source virtual disk.

CAUTION: Before using the Snapshot Virtual Disks Premium Feature in a Windows Clustered configuration, you must map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.

Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the snapshot enabling process is completed can result in the operating system misidentifying the snapshot virtual disk. This can result in data loss or an inaccessible snapshot.

Virtual Disk Copy

Virtual disk copy is a premium feature you can use to:

- Back up data.
- Copy data from disk groups that use smaller-capacity physical disks to disk groups using greater capacity physical disks.
- Restore snapshot virtual disk data to the source virtual disk.

Virtual disk copy generates a full copy of data from the source virtual disk to the target virtual disk in a storage array.

- Source virtual disk — When you create a virtual disk copy, a copy pair consisting of a source virtual disk and a target virtual disk is created on the same storage array. When a virtual disk copy is started, data from the source virtual disk is copied completely to the target virtual disk.
- Target virtual disk — When you start a virtual disk copy, the target virtual disk maintains a copy of the data from the source virtual disk. You can choose whether to use an existing virtual disk or create a new virtual disk as the target virtual disk. If you choose an existing virtual disk as the target, all data on the target is overwritten. A target virtual disk can be a standard virtual disk or the source virtual disk of a failed or disabled snapshot virtual disk.

  NOTE: The target virtual disk capacity must be equal to or greater than the source virtual disk capacity.

  When you begin the disk copy process, you must define the rate at which the copy is completed. Giving the copy process top priority slightly impacts I/O performance, while giving it lowest priority makes the copy process longer to complete. You can modify the copy priority while the disk copy is in progress. For more information, see the online help.

Virtual Disk Recovery

You can use the Edit host server-to-virtual disk mappings feature to recover data from the backup virtual disk. This functionality enables you to unmap the original source virtual disk from its host server, then map the backup virtual disk to the same host server.

Ensure that you record the LUN used to provide access to the source virtual disk. You need this information when you define a host server-to-virtual disk mapping for the target (backup) virtual disk. Also, be sure to stop all I/O activity to the source virtual disk before beginning the virtual disk recovery procedure.

Using Snapshot And Virtual Disk Copy Together

You can use the Snapshot Virtual Disk or Snapshot Virtual Disk (Legacy) and Virtual Disk Copy premium features together to back up data on the same storage array, or to restore the data on the snapshot virtual disk to its original source virtual disk.

You can copy data from a virtual disk in one of the two ways:

- By taking a point-in-time snapshot of the data
- By copying the data to another virtual disk using a virtual disk copy

You can select a snapshot virtual disk as the source virtual disk for a virtual disk copy. This configuration is one of the best ways you can apply the snapshot virtual disk feature, since it enables complete backups without any impact to the storage array I/O.

You cannot use a snapshot repository virtual disk as a source virtual disk or as a target virtual disk in a virtual disk copy. If you select the source virtual disk as the target virtual disk of a virtual disk copy, you must disable all snapshot virtual disks associated with the source virtual disk.

Multi-Path Software

Multi-path software (also referred to as the failover driver) is the software resident on the host server that provides management of the redundant data path between the host server and the storage array. For the multi-path software to correctly manage a redundant path, the configuration must have redundant iSCSI connections and cabling.

The multi-path software identifies the existence of multiple paths to a virtual disk and establishes a preferred path to that disk. If any component in the preferred path fails, the multi-path software automatically re-routes I/O requests to the alternate path so that the storage array continues to operate without interruption.

NOTE: Multi-path software is available on the MD Series storage arrays resource DVD.

Preferred And Alternate Controllers And Paths

A preferred controller is a RAID controller module designated as the owner of a virtual disk or disk group. The preferred controller is automatically selected by the MD Storage Manager when a virtual disk is created. You can change the preferred RAID controller module owner of a virtual disk after it is created. If a host is connected to only one RAID controller module, the preferred owner must manually be assigned to the RAID controller module that the host can access.

Ownership of a virtual disk is moved from the preferred controller to the secondary controller (also called the alternate controller) when the preferred controller is:

- Physically removed
- Updating firmware
- Involved in an event that caused failover to the alternate controller

Paths used by the preferred RAID controller module to access either the disks or the host server are called the preferred paths; redundant paths are called the alternate paths. If a failure causes the preferred path to become inaccessible, the storage array automatically uses the alternate path to access data, and the enclosure status LED blinks amber.

Virtual Disk Ownership

The MD Storage Manager can be used to automatically build and view virtual disks. It uses optimal settings to stripe the disk group. Virtual disks are assigned to alternating RAID controller modules when they are created. This default assignation provides a simple means for load balancing the workload of the RAID controller modules.

Ownership can later be modified to balance workload according to actual usage. If virtual disk ownership is not manually balanced, it is possible for one controller to have the majority of the work, while the other controller is idle. Limit the number of virtual disks in a disk group. If multiple virtual disks are in a disk group, consider:

- The impact each virtual disk has on other virtual disks in the same disk group.
- The patterns of usage for each virtual disk.
- Different virtual disks have higher usage at different times of day.

Load Balancing

A load balance policy is used to determine which path is used to process I/O. Multiple options for setting the load balance policies let you optimize I/O performance when mixed host interfaces are configured.

You can choose one of these load balance policies to optimize I/O performance:

- Round-robin with subset — The round-robin with subset I/O load balance policy routes I/O requests, in rotation, to each available data path to the RAID controller module that owns the virtual disks. This policy treats all paths to the RAID controller module that owns the virtual disk equally for I/O activity. Paths to the secondary RAID controller module are ignored until ownership changes. The basic assumption for the round-robin policy is that the data paths are equal. With mixed host support, the data paths may have different bandwidths or different data transfer speeds.
- Least queue depth with subset — The least queue depth with subset policy is also known as the least I/Os or least requests policy. This policy routes the next I/O request to a data path that has the least outstanding I/O requests queued. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is one of the paths in the path group of the RAID controller module that owns the virtual disk.
- Least path weight with subset (Windows operating systems only) — The least queue depth with subset policy is also known as the least I/Os or least requests policy. This policy routes the next I/O request to a data path that

has the least outstanding I/O requests queued. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is one of the paths in the path group of the RAID controller module that owns the virtual disk.

Monitoring System Performance

You can use the Performance Monitor to select virtual disks and RAID controller modules to monitor or to change the polling interval.

Keep these guidelines in mind when using the Performance Monitor:

- The Performance Monitor does not dynamically update its display if any configuration changes occur while the window is open. You must close the **Performance Monitor** window and reopen it for the changes to be displayed.
- Using the Performance Monitor to retrieve performance data can affect the normal storage array performance depending on the polling interval that you set.
- If the storage array you are monitoring begins in or transitions to an unresponsive state, an informational dialog is displayed. The dialog informs you that the Performance Monitor cannot poll the storage array for performance data.

You can view data for select virtual disks and RAID controller modules in tabular format or in graphical format.

- Table view — The data is presented in a tabular format. For more information on the data presented in the table, see the online help.
- Graphical View — The data is presented with a single x-axis and a single y-axis. The x-axis represents the time for which you selected to view performance data. The y-axis represents the metric you selected on the **Define Graph** window. A curve on the graph represents the value of the metric over time for a particular device. A historical graph can contain up to four curves. A real-time graph is limited to a single curve. For the real-time data graph, the data points refresh every 5 seconds. The system drops the oldest data point from the graph and adds the newest data point to the graph. For historical data graphs, the data points are static.

Monitoring Performance Using The Tabular View

To monitor the performance of the arrays using the tabular view:

1. In the AMW, select **Monitor → Health → Monitor Performance → Table View.**
   The **Performance Monitor** window is displayed.
2. Click **Settings**.
3. Select the items that you want to monitor.
   You can choose from the following:
   - RAID controller modules
   - Virtual disks
   - Storage array totals
4. Adjust **Polling interval** to set how often you want to update the performance statistics.

   NOTE: For an accurate elapsed time, do not use the **Synchronize RAID controller module Clocks** option while using the Performance Monitor.

   NOTE: Each time the polling interval elapses, the Performance Monitor queries the storage array again and updates the statistics in the table.
5. If required, click **OK** to go back to the **Performance Monitor** window.
6. Click **Start**.

Values for the selected storage arrays are displayed in the Performance Monitor data table. The table is updated at the interval specified in the **Polling Interval** setting.

NOTE: To force an immediate poll of the storage array, click **Update**.

7. Click **Stop** to stop monitoring the storage array.
8. To save the currently displayed performance statistics:
   a) Click **Save As** on the **Performance Monitor** window.
   b) Select the directory where you want to save the file.
   c) Type a file name in **File name**.
   d) Select a file type from the **Files of type** list.

      NOTE: The .perf extension is the default.

   e) Click **Save** to save the file.

Monitoring Performance Using The Graphical View

To monitor the performance of the arrays using the graphical view:

1. In the AMW, select **Monitor → Health → Monitor Performance → Graphical View.**

   The **Performance Monitor Graphical View** window is displayed.
2. Depending on how you want to view performance data, perform one of the following actions:

   – **To view historical data** – Select **Start** to start a monitoring session. Under **View Graph**, select **Historical data** and click **Define**.

      NOTE: When you select this option, the **Date** and **Time** fields are activated on the **Define Graph** window.

   – **To view real-time data** – Select **Real-time data**, and click **Define**.

   The **Graphical Performance Monitor** window is displayed.
3. In the **Time Interval** area, set the range of time to identify how much historical data you want to view in graphical format. Perform one of the following actions:

   – To use the first and last data point in the history buffer as the start and end time period – Select **Default**.
   – To specify the start and end date and time:

      1. Select **Customize**.
      2. In the **Date** list, select the start date and the end date of the period for which you want to view performance data. The maximum length of a history data window is seven days.
      3. In the **Time** list, select the start time and the end time for each day during the data collection period for which you want to view performance data.

         NOTE: The system uses the values you specify for the time interval to create the historical data graph.
4. In the **Graph** list, select the number of graphs you want to create.

   NOTE: You can plot five graphs at one time. You can combine devices and metrics, but you can only create and display a maximum of five graphs. The system automatically advances the graph count after the a metric/device pair are added to the **Selected table display** area (one device for real-time data, four devices for historical data).
5. In the **Metric** list, select the metric on which you want to report.

   – **Total IOs** — The number of total I/Os performed by this storage array since **Start** was selected.
   – **IO/second** — The current I/Os per second for the current polling interval.

- **MB/second** — The current MB per second, or transfer rate for the current polling interval.
- **Read Percentage** — The percentage of total I/Os that are read operations for this storage array. Write percentage can be calculated as 100 minus this value.
- **Cache Hit Percentage** — The percentage of total I/Os that are fulfilled by data from the cache rather than requiring an actual read from disk.

Values are displayed for the selected storage arrays in the Performance Monitor data table. The table is updated at the interval specified in the **Polling Interval** setting.

6. In the **Device** list, select the device you want to monitor for the metric you selected in step 5.

   You can view either real-time or historical performance metrics for the following devices:

   - RAID controller module
   - Virtual disk
   - Disk group
   - Storage array totals

7. Click **Add** to move the metric/device pair you selected to the **Selected table display** area.

   For historical data, you can select up to four devices and one metric to view per graph. For real-time data, you can select only one metric and one device to view per graph.

8. After you finish selecting the metric/device pairs on which you want to report, click **View Graph**.

   The **View Graph** window is displayed.

9. On the **View Graph** window, review the performance data displayed for the selected metric/device pair.
10. If you define more than two graphs and want to view a different graph:
    a) Click **Close**.
    b) In the **Selected** table, select a different graph, and then click **View Graph**.
11. For historical data graphs:
    a) Click **Save As** on the **Graphical Performance Monitor** window to save the currently displayed performance statistics.
    b) Select the directory where you want to save the file.
    c) Type a file name in **File name**.
    d) Select the **Comma Delimited Format file type** if you want to save the data in a form that can be imported into a commercial spreadsheet application.
    e) Click **Save** to save the file.

3
Discovering And Managing Your Storage Array

You can manage a storage array in two ways:

- Out-of-band management
- In-band management

NOTE: If you are using an iSCSI MD storage array, see *Troubleshooting the Dell PowerVault MD-Series Storage Array: Configuring the Ethernet Management Port* at **dell.com/support/manuals** for more information on the management features.

Out-Of-Band Management

In the out-of-band management method, data is separate from commands and events. Data travels through the host-to-controller interface, while commands and events travel through the management port Ethernet cables.

This management method lets you configure the maximum number of virtual disks that are supported by your operating system and host adapters.

A maximum of eight storage management stations can concurrently monitor an out-of-band managed storage array. This limit does not apply to systems that manage the storage array through the in-band management method.

When you use out-of-band management, you must set the network configuration for each RAID controller module's management Ethernet port. This includes the Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway. If you are using a Dynamic Host Configuration Protocol (DHCP) server, you can enable automatic network configuration, but if you are not using a DHCP server, you must enter the network configuration manually.

NOTE: RAID controller module network configurations can be assigned using a DHCP server (the default setting). However, if a DHCP server is not available for 150 seconds, the RAID controller modules assign static IP addresses. By default, the addresses assigned are 192.168.128.101 for controller 0 and 192.168.128.102 for controller 1.

In-Band Management

Using in-band-management, commands, events, and data travel through the host-to-controller interface. Unlike out-of-band management, commands and events are mixed with data.

NOTE: For detailed information on setting up in-band and out-of-band management see your system's Deployment Guide at **dell.com/support/manuals**.

When you add storage arrays by using this management method, specify only the host name or IP address of the host. After you add the specific host name or IP address, the host-agent software automatically detects any storage arrays that are connected to that host.

NOTE: Some operating systems can be used only as storage management stations. For more information about the operating system that you are using, see the *MD PowerVault Support Matrix* at **dell.com/support/manuals**.

For more information, see the online help topics.

Access Virtual Disk

Each RAID controller module in an MD Series storage array maintains a special virtual disk, called the access virtual disk. The host-agent software uses the access virtual disk to communicate management requests and event information between the storage management station and the RAID controller module in an in-band-managed storage array and cannot be removed without deleting the entire virtual disk, virtual disk group or virtual disk pair. The access virtual disk is not available for application data storage and cannot be removed without deleting the entire virtual disk, virtual disk group, or virtual disk pair. The default LUN is 31.

Storage Arrays

You must add the storage arrays to the MD Storage Manager before you can set up the storage array for optimal use.

NOTE: You can add storage arrays only in the EMW.

You can:

- Automatically discover storage arrays.
- Manually add storage arrays.

NOTE: Verify that your host or management station network configuration— including station IP address, subnet mask, and default gateway—is correct before adding a new storage array using the Automatic option.

NOTE: For Linux, set the default gateway so that broadcast packets are sent to 255.255.255.0. For Red Hat Enterprise Linux, if no gateway exists on the network, set the default gateway to the IP address of the NIC.

NOTE: The MD Storage Manager uses TCP/UDP port 2463 for communication to the MD storage array.

Automatic Discovery Of Storage Arrays

The Automatic Discovery process sends out a broadcast message across the local subnet and adds any storage array that responds to the message. The Automatic Discovery process finds both in-band and out-of-band storage arrays.

NOTE: The **Automatic Discovery** option and the **Rescan Hosts** option in the EMW provide automatic methods for discovering managed storage arrays.

Manual Addition Of A Storage Array

Use Manual addition if the storage array resides outside of the local subnet. This process requires specific identification information to manually add a storage array.

To add a storage array that uses out-of-band management, specify the host name or management port IP address of each controller in the storage array.

To add an in-band storage array, add the host through which the storage array is attached to the network.

NOTE: It can take several minutes for the MD Storage Manager to connect to the specified storage array.

To add a storage array manually:

1. In the EMW, select **Edit → Add Storage Array.**
2. Select the relevant management method:

- **Out-of-band management** — Enter a DNS/Network name, IPv4 address, or IPv6 address for the **RAID Controller Module** in the storage array.
- **In-band management** — Enter a name or a DNS/Network name, IPv4 address, or IPv6 address for the **Host** through which the storage array is attached to the network.

  NOTE: When adding a storage array using in-band management with iSCSI, a session must first be established between the initiator on the host server and the storage array. For more information, see Using iSCSI.

  NOTE: The host agent must be restarted before in-band management communication can be established. See Starting Or Restarting The Host Context Agent Software.

3. Click **Add**.
4. Use one of these methods to name a storage array:
   - In the EMW, select the **Setup** tab, and select **Name/Rename Storage Arrays**.
   - In the AMW, select the **Setup** tab, and select **Rename Storage Array**.
   - In the EMW, right-click the icon corresponding to the array and select **Rename**.

Setting Up Your Storage Array

A list of initial setup tasks is displayed on the **Setup** tab in the AMW. Using the tasks outlined in the **Initial Setup Tasks** area, ensures that the basic setup steps are completed.

Use the **Initial Setup Tasks** list the first time that you set up a storage array and perform the following tasks:

- Locate the storage array — Find the physical location of the storage array on your network by turning on the system identification indicator.
- Give a new name to the storage array — Use a unique name that identifies each storage array.
- Set a storage array password — Configure the storage array with a password to protect it from unauthorized access. The MD Storage Manager prompts for the password when an attempt is made to change the storage array configuration, such as when a virtual disk is created or deleted.
- Configure iSCSI host ports — Configure network parameters for each iSCSI host port automatically or specify the configuration information for each iSCSI host port.
- Configure the storage array — Create disk groups, virtual disks, and hot spare physical disks by using the Automatic configuration method or the Manual configuration method. For more information, see the online help topics.
- Map virtual disks — Map virtual disks to hosts or host groups.
- Save configuration — Save the configuration parameters in a file that you can use to restore the configuration, or reuse the configuration on another storage array. For more information, see the online help topics.

After you complete the basic steps for configuring the storage array, you can perform these optional tasks:

- Manually define hosts — Define the hosts and the host port identifiers that are connected to the storage array. Use this option only if the host is not automatically recognized and shown in the **Host Mappings** tab.
- Configure Ethernet management ports — Configure the network parameters for the Ethernet management ports on the RAID controller modules if you are managing the storage array by using the out-of-band management connections.
- View and enable premium features — Your MD Storage Manager may include premium features. View the premium features that are available and the premium features that are already started. You can start available premium features that are currently stopped.
- Manage iSCSI settings — You can configure iSCSI settings for authentication, identification, and discovery.

Locating Storage Arrays

You can use the **Blink** option to physically locate and identify a storage array. To locate the storage array:

1. Select the relevant storage array and do one of the following:
   - In the EMW, right-click the appropriate storage array, and select **Blink Storage Array**.
   - In the AMW, select the **Setup** tab, and click **Blink Storage Array**.
   - In the AMW, select **Hardware → Blink → Storage Array.**

   The LEDs on the physical disks in the storage array blink.
2. After locating the storage array, click **OK**.
   The LEDs stop blinking.
3. If the LEDs do not stop blinking, select **Hardware → Blink → Stop All Indications.**

Naming Or Renaming Storage Arrays

You can name, rename, and add comments to a storage array to facilitate identification of the storage array.

Follow these guidelines to name a storage array:

- Each storage array must be assigned a unique alphanumeric name up to 30 characters long.
- A name can consist of letters, numbers, and the special characters underscore (_), dash (–), and pound sign (#). No other special characters are allowed.

To rename a selected storage array:

1. Perform one of these actions:
   - In the AMW, select **Setup → Rename Storage Array.**
   - In the EMW, select **Devices** tab Tree view, select **Edit → Rename.**
   - In the EMW, **Devices** tab Tree view, right-click the desired array icon and select **Rename.**

   The **Rename Storage Array** dialog is displayed.
2. Type the new name of the storage array.

   NOTE: Avoid arbitrary names or names that may lose meaning in the future.
3. Click **OK**.
   A message is displayed warning you about the implications of changing the storage array name.
4. Click **Yes**.
   The new storage array name is displayed in the EMW.
5. Repeat step 1 through step 4 to name or rename additional storage arrays.

Setting A Password

You can configure each storage array with a password to protect it from unauthorized access. The MD Storage Manager prompts for the password when an attempt is made to change the storage array configuration, such as, when a virtual disk is created or deleted. View operations do not change the storage array configuration and do not require a password. You can create a new password or change an existing password.

To set a new password or change an existing password:

1. In the EMW, select the relevant storage array and open the AMW for that storage array.
   The AMW for the selected storage array is displayed.
2. In the AMW, select the **Setup** tab, and click **Set a Storage Array Password**.
   The **Set Password** dialog is displayed.

3. If you are resetting the password, type the **Current password**.

   NOTE: If you are setting the password for the first time, leave the **Current password** blank.

4. Type the **New password**.

   NOTE: It is recommended that you use a long password with at least 15 alphanumeric characters to increase security. For more information on secure passwords, see Password Guidelines.

5. Re-type the new password in **Confirm new password**.
6. Click **OK**.

   NOTE: You are not prompted for a password when you attempt to change the storage array configuration in the current management session.

Password Guidelines

- Use secure passwords for your storage array. A password should be easy for you to remember but difficult for others to determine. Consider using numbers or special characters in the place of letters, such as a 1 in the place of the letter I, or the at sign (@) in the place of the letter 'a'.
- For increased protection, use a long password with at least 15 alphanumeric characters. The maximum password length is 30 characters.
- Passwords are case sensitive.

  NOTE: You can attempt to enter a password up to ten times before the storage array enters a lockout state. Before you can try to enter a password again, you must wait 10 minutes for the storage array to reset. To reset the password, press the password reset switch on your RAID controller module.

Adding Or Editing A Comment To An Existing Storage Array

A descriptive comment, with an applicable storage array name, is a helpful identification tool. You can add or edit a comment for a storage array in the EMW only.

To add or edit a comment:

1. In the EMW, select the **Devices** tab and select the relevant managed storage array.
2. Select **Edit → Comment.**

   The **Edit Comment** dialog is displayed.

3. Type a comment.

   NOTE: The number of characters in the comment must not exceed 60 characters.

4. Click **OK**.

   This option updates the comment in the Table view and saves it in your local storage management station file system. The comment does not appear to administrators who are using other storage management stations.

Removing Storage Arrays

You can remove a storage array from the list of managed arrays if you no longer want to manage it from a specific storage management station. Removing a storage array does not affect the storage array or its data in any way. Removing a storage array only removes it from the list of storage arrays displayed in the **Devices** tab of the EMW.

To remove the storage array:

1. In the EMW, select the **Devices** tab and select the relevant managed storage array.
2. Select **Edit → Remove → Storage Array.**

   You can also right-click on a storage array and select **Remove → Storage Array.**

A message prompts you to confirm if the selected storage array is to be removed.

3. Click **Yes**.

   The storage array is removed from the list.

Enabling Premium Features

You can enable premium features on the storage array. To enable the premium features, you must obtain a feature key file specific to the premium feature that you want to enable from your storage supplier.

To enable premium features:

1. From the menu bar in the AMW, select **Storage Array → Premium Features.**

   The **Premium Features and Feature Pack Information** window is displayed.

2. Select the relevant premium feature, and click **Enable**.

   The **Select Feature Key File** dialog is displayed.

3. Navigate to the relevant folder, select the appropriate key file, and click **OK**.
4. Click **Close**.

   For more information, see the online help topics.

Displaying Failover Alert

You can change the failover alert delay for a storage array. The failover alert delay lets you delay the logging of a critical event if the multi-path driver transfers virtual disks to the non-preferred controller. If the multi-path driver transfers the virtual disks back to the preferred controller within the specified delay period, a critical event is not logged. If the transfer exceeds this delay period, then a virtual disk-not-on-preferred-path alert is issued as a critical event. You can also use this option to minimize multiple alerts when more than one virtual disk fails over because of a system error, such as a failed host adapter. For more information, see the online help topics.

To configure a failover alert delay:

1. In the AMW, on the menu bar, select **Storage Array → Change → Failover Alert Delay.**

   The **Failover Alert Delay** window is displayed.

2. In **Failover alert delay**, enter a value between 0 and 60 minutes.
3. Click **OK**.
4. If you have set a password for the selected storage array, the **Enter Password** dialog is displayed. Type the current password for the storage array.

Changing The Cache Settings On The Storage Array

To change the storage array cache settings:

1. In the AMW, select **Storage Array → Change → Cache Settings.**

   The **Change Cache Settings** window is displayed.

2. In **Start flushing**, select or enter the percentage of unwritten data in the cache to trigger a cache flush .
3. In **Stop flushing**, select or enter the percentage of unwritten data in the cache to stop a cache flush in progress.
4. Select the appropriate **Cache block size**.

   A smaller cache size is a good choice for file-system use or database-application use. A larger cache size is a good choice for applications that generate sequential I/O, such as multimedia.

5. If you have set a password for the selected storage array, the **Enter Password** dialog is displayed. Type the current password for the storage array and click **OK**.

Changing Expansion Enclosure ID Numbers

When an MD3060e Series expansion enclosure is connected to an MD Series storage array for the first time, an enclosure ID number is assigned and maintained by the expansion enclosure. This enclosure ID number is also shown in the MD Storage Manager and can be changed if required.

To change the enclosure ID numbers:

1. In the AMW, from the menu bar, select **Hardware → Enclosure → Change → ID.**
2. Select a new enclosure ID number from the **Change Enclosure ID** list.

   The enclosure ID must be between 0 and 99 (inclusive).
3. To save the changed enclosure ID, click **OK**.

Changing The Enclosure Order

You can change the order of the RAID controller modules and the expansion enclosures to match the hardware configuration in your storage array. The enclosure order change remains in effect until it is modified again.

To change the enclosure order:

1. In the AMW, from the menu bar, select **Hardware → Enclosure → Change → Hardware View Order.**
2. From the enclosures list, select the enclosure you want to move and click either **Up** or **Down** to move the enclosure to the new position.
3. Click **OK**.
4. If you have set a password for the selected storage array, the **Enter Password** dialog is displayed. Type the current password for the storage array.
5. Click **OK**.

Configuring Alert Notifications

The MD Storage Manager can send an alert for any condition on the storage array that requires your attention. Alerts can be sent as e-mail messages or as Simple Network Management Protocol (SNMP) trap messages. You can configure alert notifications either for all the storage arrays or a single storage array.

To configure alert notifications:

1. For all storage arrays, in the EMW:
   a) Select the **Setup** tab.
   b) Select **Configure Alerts**.
   c) Select **All storage arrays**.
   d) Click **OK**.

   The **Configure Alerts** dialog is displayed.
2. For a single storage array:
   a) Select the **Devices** tab.
   b) Select the relevant storage array, then select **Edit → Configure Alerts.**

   The **Configure Alerts** dialog is displayed.
3. Configure e-mail or SNMP alerts.

   For more information, see Configuring E-mail Alerts or Configuring SNMP Alerts.

Configuring E-mail Alerts

1. Open the **Configure Alerts** dialog by performing one of these actions in the EMW:
   - On the **Devices** tab, select a node and then on the menu bar, select **Edit** → **Configure Alerts.** Go to step 3.

     NOTE: This option enables you to set up alerts for all the storage arrays connected to the host.
   - On the **Setup**, select **Configure Alerts**. Go to step 2.
2. Select one of the following radio buttons to specify an alert level:
   - **All storage arrays** — Select this option to send an e-mail alert about events on all storage arrays.
   - **An individual storage array** — Select this option to send an e-mail alert about events that occur on only a specified storage array.

   These results occur, depending on your selection:
   - If you select **All storage arrays**, the **Configure Alerts** dialog is displayed.
   - If you select **An individual storage array**, the **Select Storage Array** dialog is displayed. Select the storage array for which you want to receive e-mail alerts and click **OK**. The **Configure Alerts** dialog is displayed.
   - If you do not know location of the selected storage array, click **Blink** to turn on the LEDs of the storage array.
3. In the **Configure Alerts** dialog, select the **Mail Server** tab and do the following:
   a) Type the name of the Simple Mail Transfer Protocol (SMTP) mail server.

      The SMTP mail server is the name of the mail server that forwards the e-mail alert to the configured e-mail addresses.
   b) In **Email sender address**, type the e-mail address of the sender. Use a valid e-mail address.

      The e-mail address of the sender (the network administrator) is displayed on each e-mail alert sent to the destination.
   c) (Optional) To include the contact information of the sender in the e-mail alert, select **Include contact information with the alerts**, and type the contact information.
4. Select the **Email** tab to configure the e-mail destinations:
   - Adding an e-mail address — In **Email address**, type the e-mail address, and click **Add**.
   - Replacing an e-mail address — In the **Configured email addresses** area, select the e-mail address to be replaced, type the replacement e-mail address in **Email address**, and click **Replace**.
   - Deleting an e-mail address — In the **Configured email addresses** area, select the e-mail address, and click **Delete**.
   - Validating an e-mail address — Type the e-mail address in **Email address** or select the e-mail address in the **Configured email addresses** area, and click **Test**. A test e-mail is sent to the selected e-mail address. A dialog with the results of the test and any error is displayed.

   The newly added e-mail address is displayed in the **Configured e-mail addresses** area.
5. For the selected e-mail address in the **Configured e-mail addresses** area, in the **Information To Send** list, select:
   - **Event Only** — The e-mail alert contains only the event information. By default, **Event Only** is selected.
   - **Event + Profile** — The e-mail alert contains the event information and the storage array profile.
   - **Event + Support** — The e-mail alert contains the event information and a compressed file that contains complete support information for the storage array that has generated the alert.
6. For the selected e-mail address in the **Configured e-mail addresses** area, in the **Frequency** list, select:
   - **Every event** — Sends an e-mail alert whenever an event occurs. By default, **Every event** is selected.

- **Every x hours** — Sends an e-mail alert after the specified time interval if an event has occurred during that time interval. You can select this option only if you have selected either **Event + Profile** or **Event + Support** in the **Information To Send** list.

7. Click **OK**.

   An alert icon is displayed next to each node in the Tree view where an alert is set.

8. If required, verify if the e-mail is sent successfully:

   - Provide an SMTP mail server name and an e-mail sender address for the e-mail addresses to work.
   - Ensure that the e-mail addresses that you had previously configured appear in the **Configured e-mail addresses** area.
   - Use fully qualified e-mail addresses; for example, name@mycompany.com.
   - Configure multiple e-mail addresses before you click **OK**.

Configuring SNMP Alerts

1. Open the **Configure Alerts** dialog by performing one of these actions in the EMW:

   - On the **Devices** tab, select a node and then on the menu bar, select **Edit → Configure Alerts.** Go to step 3.

     NOTE: This option enables you to set up alerts for all the storage arrays connected to the host.

   - On the **Setup**, select **Configure Alerts**. Go to step 2.

2. Select one of the following options to specify an alert level:

   - **All storage arrays** — Select this option to send an alert notification about events on all storage arrays.
   - **An individual storage array** — Select this option to send an alert notification about events that occur in only a specified storage array.

   These results occur, depending on your selection:

   - If you selected **All storage arrays**, the **Configure Alerts** dialog is displayed.
   - If you selected **An individual storage array**, the **Select Storage Array** dialog is displayed. Select the storage array for which you want to receive alert notifications and click **OK**. The **Configure Alerts** dialog is displayed.

   NOTE: If you do not know location of the selected storage array, click **Blink** to turn on the LEDs of the storage array.

3. Select the **SNMP** tab to configure the SNMP alert destinations.

   - Adding an SNMP address — In **Community name**, type the community name. In **Trap destination**, type the trap destination, and click **Add**.

     NOTE: The community name is an ASCII string that identifies a known set of network management stations and is set by the network administrator. The default community name is the string `public`. The trap destination is the IP address or the host name of a computer running an SNMP management application. An example of an SNMP enabled management application is the Dell Management Console. For more information on the Dell Management Console, see **dell.com** .

   - Replacing an SNMP address — Select the SNMP address in the **Configured SNMP addresses** area, type the replacement community name in **Community name** and the trap destination in **Trap** destination, and click **Replace**.
   - Deleting an SNMP address — Select the SNMP address in the **Configured SNMP addresses** area, and click **Delete**.
   - Validating an SNMP address — Select the SNMP address in the **Configured SNMP addresses** area, and click **Test**. A test message is sent to the SNMP address. A message box with the results of the validation and any error information is displayed.

4. Click **OK**.

An alert icon is displayed next to each node in the Tree view for which an alert is set.

NOTE:

- Any SNMP addresses that you had previously configured appear in the **Configured SNMP addresses** area.
- The **SNMP Community Name** is determined by the system administrator and configured within the a management application, such as the Dell Management Console. More information about the Dell Management Console is available at **dell.com**.
- You can configure multiple SNMP addresses before you click **OK**.

Battery Settings

A smart battery backup unit (BBU) can perform a learn cycle. The smart BBU module includes the battery, a battery gas gauge, and a battery charger. The learn cycle calibrates the smart battery gas gauge so that it provides a measurement of the charge of the battery module. A learn cycle can only start when the battery is fully charged.

The learn cycle completes the following operations:

- Discharges the battery to a predetermined threshold
- Charges the battery back to full capacity

A learn cycle starts automatically when you install a new battery module. Learn cycles for batteries in both RAID controller modules in a duplex system occur simultaneously.

Learn cycles are scheduled to start automatically at regular intervals, at the same time and on the same day of the week. The interval between cycles is described in weeks.

Use the following guidelines to adjust the interval:

- You can use the default interval.
- You can run a learn cycle at any time.
- You can set the learn cycle earlier than the currently scheduled time.
- You cannot set the learn cycle to start more than seven days later than the currently scheduled time.

Changing The Battery Settings

To change the battery settings:

1. In the AMW, from the menu bar, select **Hardware → Enclosure → Change → Battery Settings.**
   The **Battery Settings** dialog is displayed.
2. In **Battery location**, select a battery
3. Check these details about the battery:
   - Battery status
   - Battery age
   - Days until replacement

   For more information, see the online help topics.

Setting The Storage Array RAID Controller Module Clocks

You can use the **Synchronize Clocks** option to synchronize the storage array RAID controller module clocks with the storage management station. This option makes sure that the event timestamps written by the RAID controller modules

to the Event Log match the event timestamps written to host log files. The RAID controller modules remain available during synchronization.

To synchronize the RAID controller module clocks with the storage management station:

1. In the AMW, on the menu bar, select **Hardware → RAID Controller Module → Synchronize Clocks.**
2. If a password is set, in the **Enter Password** dialog, type the current password for the storage array, and click **Synchronize**.

   The RAID controller module clocks are synchronized with the management station.

4

Using iSCSI

NOTE: The following sections are relevant only to MD*xx*0i storage arrays that use the iSCSI protocol.

Changing The iSCSI Target Authentication

To change the iSCSI target authentication:

1. In the AMW, select the **Setup** tab.
2. Select **Manage iSCSI Settings**.

   The **Manage iSCSI Settings** window is displayed and by default, the **Target Authentication** tab is selected.
3. To change the authentication settings, select:

   – None — If you do not require initiator authentication. If you select **None**, any initiator can access the target.

   – CHAP — To enable an initiator that tries to authenticate the target using Challenge Handshake Authentication Protocol (CHAP). Define the CHAP secret only if you want to use mutual CHAP authentication. If you select **CHAP**, but no CHAP target secret is defined, an error message is displayed. See Creating CHAP Secrets.
4. To enter the CHAP secret, click **CHAP secret**.

   The **Enter Target CHAP Secret** dialog is displayed.
5. Enter the **Target CHAP** secret.

   The Target CHAP secret must be at least 12 characters and up to 57 characters.
6. Enter the exact target CHAP secret in **Confirm target CHAP secret**.

   NOTE: If you do not want to create a CHAP secret, you can generate a random CHAP secret automatically. To generate a random CHAP secret, click **Generate Random CHAP Secret**.
7. Click **OK**.

   NOTE: You can select the **None** and **CHAP** at the same time, for example, when one initiator may not have CHAP and the other initiator has only CHAP selected.

Entering Mutual Authentication Permissions

Mutual authentication or two-way authentication is a way for a client or a user to verify themselves to a host server, and for the host server to validate itself to the user. This validation is accomplished in such a way that both parties are sure of the other's identity.

To add mutual authentication permissions:

1. In the AMW, select the **Setup** tab.
2. Select **Manage iSCSI Settings**.

   The **Manage iSCSI Settings** window is displayed.
3. Select the **Remote Initiator Configuration** tab.
4. Select an initiator in the **Select an Initiator** area.

The initiator details are displayed.

5. Click **CHAP Secret** to enter the initiator CHAP permissions in the dialog that is displayed.
6. Click **OK.**
7. Click **OK** in the **Manage iSCSI Settings** window.

   For more information, see the online help topics.

Creating CHAP Secrets

When you set up an authentication method, you can choose to create a CHAP secret. The CHAP secret is a password that is recognized by the initiator and the target. If you are using mutual authentication to configure the storage array, you must enter the same CHAP secret that is defined in the host server iSCSI initiator, and you must define a CHAP secret on the target (the storage array) that must be configured in every iSCSI initiator that connects to the target storage array. For more information on CHAP, see Understanding CHAP Authentication in the storage array's Deployment Guide.

Initiator CHAP Secret

The initiator CHAP secret is set on the host using the iSCSI initiator configuration program provided with the host operating system. If you are using the mutual authentication method, you must define the initiator CHAP secret when you set up the host. This must be the same CHAP secret that is defined for the target when defining mutual authentication settings.

Target CHAP Secret

If you are using CHAP secrets, you must define the CHAP secret for the target.

Valid Characters For CHAP Secrets

The CHAP secret must be between 12 and 57 characters. The CHAP secret supports characters with ASCII values of 32 to 126 decimal. See the following table for a list of valid ASCII characters.

Space	!	“	#	$	%	&	'	(	)	*	+
,	-	.	/	0	1	2	3	4	5	6	7
8	9	:	;	<	=	>	?	@	A	B	C
D	E	F	G	H	I	J	K	L	M	N	O
P	Q	R	S	T	U	V	W	X	Y	Z	[
\	]	^	_	a	b	c	d	e	f	g	h
I	j	k	l	m	n	o	p	q	r	s	t
u	v	w	x	y	z	{	|	}	~		

Changing The iSCSI Target Identification

You cannot change the iSCSI target name, but you can associate an alias with the target for simpler identification. Aliases are useful because the iSCSI target names are not intuitive. Provide an iSCSI target alias that is meaningful and easy to remember.

To change the iSCSI target identification:

1. In the AMW, select the **Setup** tab.
2. Select **Manage iSCSI Settings**.
   The **Manage iSCSI Settings** window is displayed.
3. Select the **Target Configuration** tab.
4. Type the alias in **iSCSI alias**.
5. Click **OK**.

   NOTE: Aliases can contain up to 30 characters. Aliases can include letters, numbers, and the special characters underscore (_), minus (-), and pound sign (#). No other special characters are permitted.

   NOTE: Open iSCSI (which is used by Red Hat Enterprise Linux 5 and SUSE Linux Enterprise Server 10 with SP 1) does not support using target alias.

Changing The iSCSI Target Discovery Settings

To change the iSCSI target discovery settings:

1. In the AMW, select the **Setup** tab.
2. Select **Manage iSCSI Settings**.
   The **Manage iSCSI Settings** window is displayed.
3. Select the **Target Discovery** tab.
4. Select **Use iSNS** to activate iSCSI target discovery.
5. To activate iSCSI target discovery, you can use one of the following methods:
   - Select **Obtain configuration automatically from DHCP server** to automatically activate target discovery for IPv4 settings using the Dynamic Host Configuration Protocol (DHCP). You can also refresh the DHCP.
   - Select **Specify Configuration**, and type the IPv4 address to activate the target discovery.
   - Type the **iSNS server IP address** in the IPv6 settings area to activate the target discovery.

   NOTE: After you manually enter an IP address, you can also click **Advanced** to configure the customized TCP listening ports.

   NOTE: If you do not want to allow discovery sessions that are not named, select **Disallow un-named discovery sessions**.

   NOTE: Un-named discovery sessions are discovery sessions that are permitted to run without a target name. With an un-named discovery session, the target name or the target portal group tag is not available to enforce the iSCSI session identifier (ISID) rule.
6. Click **OK**.

Configuring The iSCSI Host Ports

The default method for configuring the iSCSI host ports, for IPv4 addressing, is DHCP. Always use this method unless your network does not have a DHCP server. It is advisable to assign static DHCP addresses to the iSCSI ports to ensure

continuous connectivity. For IPv6 addressing, the default method is Stateless auto-configuration. Always use this method for IPv6. For more information, see *Troubleshooting your Dell PowerVault MD-Series Storaeg Array: Configuring iSCSI Data Ports* at **dell.com/support/manuals**.

To configure the iSCSI host ports:

1. In the AMW, select the **Setup** tab.
2. Select **Configure iSCSI Host Ports**.

   The **Configure iSCSI Ports** window is displayed.
3. In the **iSCSI port** list, select an appropriate RAID controller module and an iSCSI host port.

   The connection status between the storage array and the host is displayed in the Status area when you select an iSCSI host port. The connection status is either connected or disconnected. Additionally, the media access control address (MAC) of the selected iSCSI host port is displayed in the MAC address area.

   NOTE: For each iSCSI host port, you can use either IPv4 settings or IPv6 settings or both.
4. In the **Configured Ethernet port speed** list, select a network speed for the iSCSI host port.

   The network speed values in the **Configured Ethernet port speed** list depend on the maximum speed that the network can support. Only the network speeds that are supported are displayed.

   All of the host ports on a single controller operate at the same speed. An error is displayed if different speeds are selected for the host ports on the same controller.
5. To use the IPv4 settings for the iSCSI host port, select **Enable IPv4** and select the **IPv4 Settings** tab.
6. To use the IPv6 settings for the iSCSI host port, select **Enable IPv6** and select the **IPv6 Settings** tab.
7. To configure the IPv4 and IPv6 settings, select:

   - **Obtain configuration automatically from DHCP server** to automatically configure the settings. This option is selected by default.
   - **Specify configuration** to manually configure the settings.

   NOTE: If you select the automatic configuration method, the configuration is obtained automatically using the DHCP for IPv4 settings. Similarly for IPv6 settings, the configuration is obtained automatically based on the MAC address and the IPv6 routers present on the subnetwork.
8. Click **Advanced IPv4 Settings** and **Advanced IPv6 Settings** to configure the Virtual Local Area Network (VLAN) support and Ethernet priority.
9. Click the **Advanced Port Settings** to configure the **TCP listening port settings** and **Jumbo frame** settings.
10. To enable the Internet Control Message Protocol (ICMP), select Enable **ICMP PING responses**.

    The ICMP setting applies to all the iSCSI host ports in the storage array configured for IPv4 addressing.

    NOTE: The ICMP is one of the core protocols of the Internet Protocol suite. The ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.
11. Click **OK**.

Advanced iSCSI Host Port Settings

NOTE: Configuring the advanced iSCSI host ports settings is optional.

Use the advanced settings for the individual iSCSI host ports to specify the TCP frame size, the virtual LAN, and the network priority.

Setting	Description
Virtual LAN (VLAN)	A method of creating independent logical networks within a physical network. Several VLANs can exist within a network. VLAN 1 is the default VLAN.

Setting	Description
	NOTE: For more information on creating and configuring a VLAN with MD Support Manager, in the AMW, click the **Support** tab, then click **View Online Help**.
Ethernet Priority	The network priority can be set from lowest to highest. Although network managers must determine these mappings, the IEEE has made broad recommendations: • 0 — lowest priority (default). • (1–4) — ranges from "loss eligible" traffic to controlled-load applications, such as streaming multimedia and business-critical traffic. • (5–6) — delay-sensitive applications such as interactive video and voice. • 7 — highest priority reserved for network-critical traffic.
TCP Listening Port	The default Transmission Control Protocol (TCP) listening port is 3260.
Jumbo Frames	The maximum transmission units (MTUs). It can be set between 1501 and 9000 Bytes per frame. If the Jumbo Frames are disabled, the default MTU is 1500 Bytes per frame.

NOTE: Changing any of these settings resets the iSCSI port. I/O is interrupted to any host accessing that port. You can access the I/O automatically after the port restarts and the host logs in again.

Viewing Or Ending An iSCSI Session

You may want to end an iSCSI session for the following reasons:

- Unauthorized access — If an initiator is logged on whom you consider to not have access, you can end the iSCSI session. Ending the iSCSI session forces the initiator to log off the storage array. The initiator can log on if **None** authentication method is available.
- System downtime — If you need to turn off a storage array and initiators are logged on, you can end the iSCSI session to log off the initiators from the storage array.

To view or end an iSCSI session:

1. In the AMW menu bar, select **Storage Array** → **iSCSI** → **End Sessions.**
2. Select the iSCSI session that you want to view in the **Current sessions** area.

   The details are displayed in the **Details** area. Click **Save As** to save the entire iSCSI sessions topology as a text file.
3. To end the session:
   a) Select the session that you want to end, and then click **End Session**.

      The **End Session confirmation** window is displayed.
   a) Type `yes` to confirm that you want to end the iSCSI session.
   b) Click **OK**.

   NOTE: If you end a session, any corresponding connections terminate the link between the host and the storage array, and the data on the storage array is no longer available.

   NOTE: When a session is manually terminated using the MD Storage Manager, the iSCSI initiator software automatically attempts to re-establish the terminated connection to the storage array. This may cause an error message.

Viewing iSCSI Statistics And Setting Baseline Statistics

To view iSCSI statistics and set baseline statistics:

1. In the AMW menu bar, select **Monitor → Health → iSCSI Statistics.**
   The **View iSCSI Statistics** window is displayed.
2. Select the iSCSI statistic type you want to view in the **iSCSI Statistics Type** area. You can select:
   - Ethernet MAC statistics
   - Ethernet TCP/IP statistics
   - Target (protocol) statistics
   - Local initiator (protocol) statistics
3. In the **Options** area, select:
   - Raw statistics — To view the raw statistics. Raw statistics are all the statistics that have been gathered since the RAID controller modules were started.
   - Baseline statistics — To view the baseline statistics. Baseline statistics are point-in-time statistics that have been gathered since you set the baseline time.

   After you select the statistics type and either raw or baseline statistics, the details of the statistics appear in the statistics tables.

   NOTE: You can click **Save As** to save the statistics that you are viewing in a text file.
4. To set the baseline for the statistics:
   a) Select **Baseline statistics**.
   b) Click **Set Baseline**.
   c) Confirm that you want to set the baseline statistics in the dialog that is displayed.

   The baseline time shows the latest time you set the baseline. The sampling interval is the difference in time from when you set the baseline until you launch the dialog or click **Refresh**.

   NOTE: You must first set a baseline before you can compare baseline statistics.

Edit, Remove, Or Rename Host Topology

If you give access to the incorrect host or the incorrect host group, you can remove or edit the host topology. Follow the appropriate procedures given in the following table to correct the host topology.

Table 2. Host Topology Actions

Desired Action	Steps to Complete Action
Move a host Move a host group	1. Click the **Host Mappings** tab. 2. Select the Host that you want to move, and then select **Host Mappings → Move.** 3. Select a host group to move the host to and click **OK**.
Manually delete the host and the host group	1. Click the **Host Mappings** tab. 2. Select the item that you want to remove and select **Host Mappings → Remove.**
Rename the host or the host group	1. Click the **Host Mappings** tab. 2. Select the item that you want to remove and select **Host Mappings → Rename.** 3. Type a new label for the host and click **OK**.

For more information about Host, Host Groups, and Host Topology, see About Your Host.

5

Event Monitor

An event monitor is provided with Dell PowerVault Modular Disk Storage Manager (MD Storage Manager). The event monitor runs continuously in the background and monitors activity on the managed storage arrays. If the event monitor detects any critical problems, it can notify a host or remote system using e-mail, Simple Network Management Protocol (SNMP) trap messages, or both.

For the most timely and continuous notification of events, enable the event monitor on a management station that runs 24 hours a day. Enabling the event monitor on multiple systems or having a combination of an event monitor and MD Storage Manager active can result in duplicate events, but this does not indicate multiple failures on the array.

The Event Monitor is a background task that runs independently of the Enterprise Management Window (EMW).

To use the Event Monitor, perform one of these actions:

- Set up alert destinations for the managed device that you want to monitor. A possible alert destination would be the Dell Management Console. For more information on the Dell Management Console, see **dell.com**.
- Replicate the alert settings from a particular managed device by copying the `emwdata.bin` file to every storage management station from which you want to receive alerts.

Each managed device shows a check mark that indicates that alerts have been set.

Enabling Or Disabling The Event Monitor

You can enable or disable the event monitor at any time.

Disable the event monitor if you do not want the system to send alert notifications. If you are running the event monitor on multiple systems, disabling the event monitor on all but one system prevents the sending of duplicate messages.

NOTE: It is recommended that you configure the event monitor to start by default on a management station that runs 24 hours a day.

Windows

To enable or disable the event monitor:

1. Click **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services** or click **Start** → **Administrative Tools** → **Services.**
2. From the list of services, select **Modular Disk Storage Manager Event Monitor**.
3. Select **Action** → **Properties.**
4. To enable the event monitor, in the **Service Status** area, click **Start**.
5. To disable the event monitor, in the **Service Status** area, click **Stop**.

Linux

To enable the event monitor, at the command prompt, type `SMmonitor start` and press <Enter>. When the program startup begins, the following message is displayed: `SMmonitor started.`

To disable the event monitor, start terminal emulation application (console ox xterm) and at the command prompt, type `SMmonitor stop`, and press <Enter>. When the program shutdown is complete, the following message is displayed: `Stopping Monitor process.`

6

About Your Host

Configuring Host Access

Dell PowerVault Modular Disk Storage Manager (MD Storage Manager) is comprised of multiple modules. One of these modules is the Host Context Agent, which is installed as part of the MD Storage Manager installation and runs continuously in the background.

If the Host Context Agent is running on a host, that host and the host ports connected from it to the storage array are automatically detected by the MD Storage Manager. The host ports are displayed in the **Host Mappings** tab in the Array Management Window (AMW). The host must be manually added under the **Default Host Group** in the **Host Mappings** tab.

NOTE: On MD3260i and MD3660i storage arrays that use the iSCSI protocol, the Host Context Agent is not dynamic and must be restarted after establishing iSCSI sessions to automatically detect them.

Use the Define Host Wizard to define the hosts that access the virtual disks in the storage array. Defining a host is one of the steps required to let the storage array know which hosts are attached to it and to allow access to the virtual disks. For more information on defining the hosts, see Defining A Host.

To enable the host to write to the storage array, you must map the host to the virtual disk. This mapping grants a host or a host group access to a particular virtual disk or to a number of virtual disks in a storage array. You can define the mappings on the **Host Mappings** tab in the AMW.

On the **Summary** tab in the AMW, the **Host Mappings** area indicates how many hosts are configured to access the storage array. Click **Configured Hosts** in the **Host Mappings** area to see the names of the hosts.

A collection of elements, such as default host groups, hosts, and host ports, are displayed as nodes in the object tree on the left pane of the **Host Mappings** tab.

The host topology is reconfigurable. You can perform the following tasks:

- Create a host and assign an alias or user label.
- Add or associate a new host port identifier to a particular host.
- Change the host port identifier alias or user label.
- Move or associate a host port identifier to a different host.
- Replace a host port identifier with a new host port identifier.
- Manually activate an inactive host port so that the port can gain access to host specific or host group specific LUN mappings.
- Set the host port type to another type.
- Move a host from one host group to another host group.
- Remove a host group, a host, or a host port identifier.
- Rename a host group, or a host.

Using The Host Mappings Tab

In the **Host Mappings** tab, you can:

- Define hosts and hosts groups

- Add mappings to the selected host groups

For more information, see the online help topics.

Defining A Host

You can use the Define Host Wizard in the AMW to define a host for a storage array. Either a known unassociated host port identifier or a new host port identifier can be added.

A user label must be specified before the host port identifier may be added (the **Add** button is disabled until one is entered).

To define a host:

1. In the AMW, select the **Host Mappings** tab.
2. Perform one of the actions:
   - From the menu bar, select **Host Mappings** → **Define** → **Host.**
   - Select the **Setup** tab, and click **Manually Define Hosts**.
   - Select the **Host Mappings** tab. Right-click the root node (storage array name), Default Group node, or Host Group node in the object tree to which you want to add the host, and select **Define** → **Host** from the pop-up menu.

   The **Specify Host Name** window is displayed.
3. In **Host name**, enter an alphanumeric name of up to 30 characters.
4. Select the relevant option in **Do you plan to use the storage partitions in this storage array?** and click **Next**.

   The **Specify Host Port Identifiers** window is displayed.
5. Select the relevant option to add a host port identifier to the host, you can select:
   - **Add by selecting a known unassociated host port identifier** — In **Known unassociated host port identifier**, select the relevant host port identifier.
   - **Add by creating a new host port identifier** — In **New host port identifier**, enter a 16 character name and an **Alias** of up to 30 characters for the host port identifier, and click **Add**.

     NOTE: The host port identifier name must contain only the letters A through F.
6. Click **Add**.

   The host port identifier and the alias for the host port identifier is added to the host port identifier table.
7. Click **Next**.

   The **Specify Host Type** window is displayed.
8. In **Host type** (operating system), select the relevant operating system for the host.

   The **Host Group Question** window is displayed.
9. In the **Host Group Question** window, you can select:
   - **Yes** — This host shares access to the same virtual disks with other hosts.
   - **No** — This host does NOT share access to the same virtual disks with other hosts.
10. Click **Next**.
11. If you select:
    - **Yes** — The **Specify Host Group** window is displayed.
    - **No** — Go to step 13.
12. Enter the name of the host group or select an existing host group and click **Next**.

    The **Preview** window is displayed.

13. Click **Finish**.
    The **Creation Successful** window is displayed confirming that the new host is created.
14. To create another host, click **Yes** on the **Creation Successful** window.

Removing Host Access

To remove host access:

1. In the AMW, select the **Host Mappings** tab.
2. Select the host node from the object tree on the left pane.
3. Perform one of these actions:
   - From the menu bar, select **Host Mappings → Host → Remove.**
   - Right-click the host node, and select **Remove** from the pop-up menu.

   The **Remove confirmation** dialog is displayed.
4. Type `yes`.
5. Click **OK**.

Managing Host Groups

A host group is a logical entity of two or more hosts that share access to specific virtual disks on the storage array. You create host groups using the MD Storage Manager.

All hosts in a host group must have the same host type (operating system). In addition, all hosts in the host group must have special software, such as clustering software, to manage virtual disk sharing and accessibility.

If a host is part of a cluster, every host in the cluster must be connected to the storage array, and every host in the cluster must be added to the host group.

Creating A Host Group

To create a host group:

1. In the AMW, select the **Host Mappings** tab.
2. In the object tree, select the storage array or the **Default Group**.
3. Perform one of the following actions:
   - From the menu bar, select **Host Mappings → Define → Host Group.**
   - Right-click the storage array or the **Default Group**, and select **Define → Host Group** from the pop-up menu.

   The **Define Host Group** window is displayed.
4. Type the name of the new host group in **Enter new host group name**.
5. Select the appropriate hosts in the **Select hosts to add** area.
6. Click **Add**.
   The new host is added in the **Hosts in group** area.

   NOTE: To remove hosts, select the hosts in the **Hosts in group** area, and click **Remove**.
7. Click **OK**.

Adding A Host To A Host Group

You can add a host to an existing host group or a new host group using the **Define Host Wizard**. For more information, see Defining A Host.

You can also move a host to a different host group. For more information, see Moving A Host To A Different Host Group.

Removing A Host From A Host Group

You can remove a host from the object tree on the **Host Mappings** tab of the AMW. For more information, see Removing A Host Group.

Moving A Host To A Different Host Group

To move a host to a different host group:

1. In the AMW, select the **Host Mappings** tab, select the host node in the object tree.
2. Perform one of these actions:
   - From the menu bar, select **Host Mappings** → **Host** → **Move.**
   - Right-click the host node, and select **Move** from the pop-up menu.

   The **Move Host** dialog is displayed.
3. In the **Select host group** list, select the host group to which you want to move the host.
   You can also move the host out of the host group and add it under the default group.
   The **Move Host** confirmation dialog is displayed.
4. Click **Yes**.
   The host is moved to the selected host group with the following mappings:
   - The host retains the specific virtual disk mappings assigned to it.
   - The host inherits the virtual disk mappings assigned to the host group to which it is moved.
   - The host loses the virtual disk mappings assigned to the host group from which it was moved.

Removing A Host Group

To remove a host group:

1. In the AMW, select the **Host Mappings** tab, select the host group node in the object tree.
2. Perform one of these actions:
   - From the menu bar, select **Host Mappings** → **Host Group** → **Remove.**
   - Right-click the host group node, and select **Remove** from the pop-up menu.

   The **Remove** dialog is displayed.
3. Click **Yes**.
   The selected host group is removed.

Host Topology

Host topology is the organization of hosts, host groups, and host interfaces configured for a storage array. You can view the host topology in the **Host Mappings** tab of the AMW. For more information, see Using The Host Mappings Tab.

The following tasks change the host topology:

- Moving a host or a host connection
- Renaming a host group, a host, or a host connection
- Adding a host connection
- Replacing a host connection
- Changing a host type

The MD Storage Manager automatically detects these changes for any host running the host agent software.

Starting Or Stopping The Host Context Agent

The Host Context Agent discovers the host topology. The Host Context Agent starts and stops with the host. The topology discovered by the Host Context Agent can be viewed by clicking **Configure Host Access (Automatic)** in the **Configure** tab in the MD Storage Manager.

You must stop and restart the Host Context Agent to see the changes to the host topology if:

- A new storage array is attached to the host server.
- A host is added while turning on power to the RAID controller modules.

To start or stop the Host Context Agent on Linux, enter the following commands at the prompt:

```
SMagent start
```

```
SMagent stop
```

You must stop and then restart SMagent after:

- Moving a controller offline or replacing a controller.
- Removing host-to-array connections from or attaching host-to-array connections to a Linux host server.

To start or stop the Host Context Agent on Windows:

1. Do one of the following:
   - Click **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**
   - Click **Start** → **Administrative Tools** → **Services**
2. From the list of services, select **Modular Disk Storage Manager Agent**.
3. If the Host Context Agent is running, click **Action** → **Stop,** then wait approximately 5 seconds.
4. Click **Action** → **Start.**

I/O Data Path Protection

You can have multiple host-to-array connections for a host. Ensure that you select all the connections to the array when configuring host access to the storage array.

NOTE: See the Deployment Guide for more information on cabling configurations.

NOTE: For more information on configuring hosts, see About Your Host.

If a component such as a RAID controller module or a cable fails, or an error occurs on the data path to the preferred RAID controller module, virtual disk ownership is moved to the alternate non preferred RAID controller module for processing. This failure or error is called failover.

Drivers for multi-path frameworks such as Microsoft Multi-Path IO (MPIO) and Linux Device Mapper (DM) are installed on host systems that access the storage array and provide I/O path failover.

For more information on Linux DM, see Device Mapper Multipath for Linux. For more information on MPIO, see **microsoft.com**.

NOTE: You must have the multi-path driver installed on the hosts at all times, even in a configuration where there is only one path to the storage system, such as a single port cluster configuration.

During a failover, the virtual disk transfer is logged as a critical event, and an alert notification is sent automatically if you have configured alert destinations for the storage array.

Managing Host Port Identifiers

You can do the following to manage the host port identifiers that are added to the storage array:

- Add — Add or associate a new host port identifier to a particular host.
- Edit — Change the host port identifier alias or user label. You can move (associate) the host port identifier to a new host.
- Replace — Replace a particular host port identifier with another host port identifier.
- Remove — Remove the association between a particular host port identifier and the associated host.

To manage a host port identifier:

1. In the AMW, select the **Host Mappings** tab.
2. Perform one of these actions:
   - Right-click the host in the object tree, and select **Manage Host Port Identifiers** in the pop-up menu.
   - From the menu bar, select **Host Mappings → Manage Host Port Identifiers.**

   The **Manage Host Port Identifiers** dialog is displayed.
3. To manage the host port identifiers in the **Show host port identifiers associated with** list:
   - For a specific host, select the host from the list of hosts that are associated with the storage array.
   - For all hosts, select **All hosts** from the list of hosts that are associated with the storage array.
4. If you are adding a new host port identifier, go to step 5. If you are managing an existing host port identifier, go to step 10.
5. Click **Add**.

   The **Add Host Port Identifier** dialog is displayed.
6. Select the appropriate host interface type.
7. Select the method to add a host port identifier to the host. You can select:
   - **Add by selecting a known unassociated host port identifier** — Select the appropriate host port identifier from the existing list of **Known unassociated host port identifiers**.
   - **Add by creating a new host port identifier** — In **New host port identifier**, enter the name of the new host port identifier.
8. In **Alias**, enter an alphanumeric name of up to 30 characters.
9. In **Associated with host**, select the appropriate host.

The newly added host port identifier is added to the **Host port identifier information** area.

10. Select the host port identifier that you want to manage from the list of host port identifiers in the **Host port identifier information** area.
11. Perform one of these actions for the selected host port identifier:

    - To edit the host port identifier — Select the appropriate host port identifier and click **Edit**. The **Edit Host Port Identifier** dialog is displayed. Update **User label** and **Associated with host** and click **Save**.
    - To replace the host port identifier — Select the appropriate host port identifier and click **Replace**. The **Replace Host Port Identifier** dialog is displayed. Replace the current host port identifier with a known unassociated host port identifier or create a new host port identifier, update **User label** and click **Replace**.
    - To remove the host port identifier — Select the appropriate host port identifier and click **Edit**. The **Remove Host Port Identifier** dialog is displayed. Type `yes` and click **OK**.

    For more information, see the online help topics.

7

Disk Groups, Standard Virtual Disks, And Thin Virtual Disks

Creating Disk Groups And Virtual Disks

Disk groups are created in the unconfigured capacity of a storage array, and virtual disks are created in the free capacity of a disk group or disk pool. The maximum number of physical disks supported in a disk group is 120 (192 with the premium feature activated). The hosts attached to the storage array read and write data to the virtual disks.

NOTE: Before you can create virtual disks, you must first organize the physical disks into disk groups and configure host access. Then you can create virtual disks within a disk group.

To create a virtual disk, use one of the following methods:

- Create a new disk group from unconfigured capacity. First define the RAID level and free capacity (available storage space) for the disk group, and then define the parameters for the first virtual disk in the new disk group.
- Create a new virtual disk in the free capacity of an existing disk group or disk pool. You only need to specify the parameters for the new virtual disk.

A disk group has a set amount of free capacity that is configured when the disk group is created. You can use that free capacity to subdivide the disk group into one or more virtual disks.

You can create disk groups and virtual disks using:

- Automatic configuration — Provides the fastest method, but with limited configuration options
- Manual configuration — Provides more configuration options

When creating a virtual disk, consider the uses for that virtual disk, and select an appropriate capacity for those uses. For example, if a disk group has a virtual disk that stores multimedia files (which tend to be large) and another virtual disk that stores text files (which tend to be small), the multimedia file virtual disk requires more capacity than the text file virtual disk.

A disk group should be organized according to its related tasks and subtasks. For example, if you create a disk group for the Accounting Department, you can create virtual disks that match the different types of accounting performed in the department: Accounts Receivable (AR), Accounts Payable (AP), internal billing, and so forth. In this scenario, the AR and AP virtual disks probably need more capacity than the internal billing virtual disk.

NOTE: In Linux, the host must be rebooted after deleting virtual disks to reset the **/dev** entries.

NOTE: Before you can use a virtual disk, you must register the disk with the host systems. See Host-To-Virtual Disk Mapping.

Creating Disk Groups

NOTE: If you have not created disk groups for a storage array, the **Disk Pool Automatic Configuration** Wizard is displayed when you open the AMW. For more information on creating storage space from disk pools, see Disk Pools.

NOTE: Thin-provisioned virtual disks can be created from disk pools. If you are not using disk pools, only standard virtual disks can be created. For more information, see Thin Virtual Disks.

You can create disk groups either using **Automatic** configuration or **Manual** configuration.

To create disk groups:

1. To start the **Create Disk Group** Wizard, perform one of these actions:
   - To create a disk group from unconfigured capacity in the storage array, in the **Storage & Copy Services** tab, select a storage array and right-click the **Total Unconfigured Capacity** node, and select **Create Disk Group** from the pop-up menu.
   - To create a disk group from unassigned physical disks in the storage array — On the **Storage & Copy Services** tab, select one or more unassigned physical disks of the same physical disk type, and from the menu bar, select **Storage → Disk Group → Create.**
   - Select the **Hardware** tab and right-click the unassigned physical disks, and select **Create Disk Group** from the pop-up menu.
   - To create a secure disk group — On the **Hardware** tab, select one or more unassigned security capable physical disks of the same physical disk type, and from the menu bar, select **Storage → Disk Group → Create.**

   The **Introduction (Create Disk Group)** window is displayed.
2. Click **Next**.
   The **Disk Group Name & Physical Disk Selection** window is displayed.
3. Type up to 30-character name of the disk group in **Disk group name**.
4. Select the appropriate **Physical Disk selection choices** and click **Next.**
   You can make the following choices:
   - **Automatic**.
   - **Manual**.
5. For automatic configuration, the **RAID Level and Capacity** window is displayed:
   a) Select the appropriate RAID level in **Select RAID level**. You can select RAID levels 0, 1/10, 5, and 6.
   Depending on your RAID level selection, the physical disks available for the selected RAID level are displayed in **Select capacity** table.
   b) In the **Select Capacity** table, select the relevant disk group capacity, and click **Finish**.
6. For manual configuration, the **Manual Physical Disk Selection** window is displayed:
   a) Select the appropriate RAID level in **Select RAID level**. You can select RAID levels 0, 1/10, 5, and 6.
   Depending on your RAID level selection, the physical disks available for the selected RAID level are displayed in **Unselected physical disks** table.
   b) In the **Unselected physical disks** table, select the appropriate physical disks and click **Add**.

   NOTE: You can select multiple physical disks at the same time by holding <Ctrl> or <Shift> and selecting additional physical disks.
   c) To view the capacity of the new disk group, click **Calculate Capacity**.
   d) Click **Finish**.

   A message prompts you that the disk group is successfully created and that you should create at least one virtual disk before you can use the capacity of the new disk group. For more information on creating virtual disks, see Creating Virtual Disks.

Locating A Disk Group

You can physically locate and identify all of the physical disks that comprise a selected disk group. An LED blinks on each physical disk in the disk group.

To locate a disk group:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Right-click on a disk group and select **Blink** from the pop-up menu.

   The LEDs for the selected disk group blink.
3. After locating the disk group, click **OK**.

   The LEDs stop blinking.
4. If the LEDs for the disk group do not stop blinking, from the toolbar in AMW, select **Hardware → Blink → Stop All Indications.**

   If the LEDs successfully stop blinking, a confirmation message is displayed.
5. Click **OK**.

Creating Standard Virtual Disks

Keep these important guidelines in mind when you create a standard virtual disk:

- Many hosts can have 256 logical unit numbers (LUNs) mapped per storage partition, but the number varies per operating system.
- After you create one or more virtual disks and assign a mapping, you must register the virtual disk with the operating system. In addition, you must make sure that the host recognizes the mapping between the physical storage array name and the virtual disk name. Depending on the operating system, run the host-based utilities, **hot_add** and **SMdevices**.
- If the storage array contains physical disks with different media types or different interface types, multiple **Unconfigured Capacity** nodes may be displayed in the **Total Unconfigured Capacity** pane of the **Storage & Copy Services** tab. Each physical disk type has an associated **Unconfigured Capacity** node if unassigned physical disks are available in the expansion enclosure.
- You cannot create a disk group and subsequent virtual disk from different physical disk technology types. Each physical disk that comprises the disk group must be of the same physical disk type.

  NOTE: Ensure that you create disk groups before creating virtual disks. If you chose an **Unconfigured Capacity** node or unassigned physical disks to create a virtual disk, the **Disk Group Required** dialog is displayed. Click **Yes** and create a disk group by using the **Create Disk Group Wizard**. The **Create Virtual Disk Wizard** is displayed after you create the disk group.

To create standard virtual disks:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select a **Free Capacity** node from an existing disk group and do one of the following:

   – From the menu bar, select **Storage → Virtual Disk → Create → Virtual Disk.**

   – Right click on the **Free Capacity** and select **Create Disk Group**.

   The **Create Virtual Disk: Specify Parameters** window is displayed.
3. Select the appropriate unit for memory in **Units** and enter the capacity of the virtual disk in **New virtual disk capacity**.
4. In **Virtual disk name**, enter a virtual disk name of up to 30 characters.
5. In the **Map to host** list, select an appropriate host or select **Map later**.

6. In the **Data Service (DS) Attributes** area, you can select:
   - **Enable data assurance (DA) protection on the new virtual disk**
   - **Use SSD cache**
7. In the **Virtual disk I/O characteristics type** list, select the appropriate Virtual Disk I/O characteristics type. You can select:
   - **File system (typical)**
   - **Database**
   - **Multimedia**
   - **Custom**

   NOTE: If you select **Custom**, you must select an appropriate segment size.
8. Select **Enable dynamic cache read prefetch**.

   For more information on virtual disk cache settings, see Changing The Virtual Disk Cache Settings.

   NOTE: Enable dynamic cache read prefetch must be disabled if the virtual disk is used for database applications or applications with a large percentage of random reads.
9. From the **Segment size** list, select an appropriate segment size.
10. Click **Finish**.

    The virtual disks are created.

    NOTE: A message prompts you to confirm if you want to create another virtual disk. Click **Yes** to proceed further, else click **No**.

    NOTE: Thin virtual disks are supported on disk pools. For more information, see Thin Virtual Disks.

Changing The Virtual Disk Modification Priority

You can specify the modification priority setting for a single virtual disk or multiple virtual disks on a storage array.

Guidelines to change the modification priority of a virtual disk:

- If more than one virtual disk is selected, the modification priority defaults to the lowest priority. The current priority is shown only if a single virtual disk is selected.
- Changing the modification priority by using this option modifies the priority for the selected virtual disks.

To change the virtual disk modification priority:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select a virtual disk.
3. In the menu bar, select **Storage → Virtual Disk → Change → Modification Priority.**

   The **Change Modification Priority** window is displayed.
4. Select one or more virtual disks. Move the Select modification priority slider bar to the desired priority.

   NOTE: To select nonadjacent virtual disks, press <Ctrl> click and select the appropriate virtual disks. To select adjacent virtual disks, press <Shift> click the appropriate virtual disks. To select all of the available virtual disks, click **Select All**.
5. Click OK.

   A message prompts you to confirm the change in the virtual disk modification priority.
6. Click **Yes**.
7. Click **OK**.

Changing The Virtual Disk Cache Settings

You can specify the cache memory settings for a single virtual disk or for multiple virtual disks in a storage array.

Guidelines to change cache settings for a virtual disk:

- After opening the **Change Cache Settings** dialog, the system may display a window indicating that the RAID controller module has temporarily suspended caching operations. This action may occur when a new battery is charging, when a RAID controller module has been removed, or if a mismatch in cache sizes has been detected by the RAID controller module. After the condition has cleared, the cache properties selected in the dialog become active. If the selected cache properties do not become active, contact your Technical Support representative.
- If you select more than one virtual disk, the cache settings default to no settings selected. The current cache settings appear only if you select a single virtual disk.
- If you change the cache settings by using this option, the priority of all of the virtual disks that you selected is modified.

To change the virtual disk cache settings:

1. In the AMW, select the Storage & Copy Services tab and select a virtual disk.
2. In the menu bar, select **Storage → Virtual Disk → Change → Cache Settings.**

   The **Change Cache Settings** window is displayed.
3. Select one or more virtual disks.

   To select nonadjacent virtual disks, press <Ctrl> click. To select adjacent virtual disks, press <Shift> click. To select all of the available virtual disks, select **Select All**.
4. In the **Cache Properties** area, you can select:
   - **Enable read caching**
   - **Enable write caching**
     - **Enable write caching without batteries** — to permit write caching to continue even if the RAID controller module batteries are discharged completely, not fully charged, or are not present.
     - **Enable write caching with mirroring** — to mirror cached data across two redundant RAID controller modules that have the same cache size.
   - **Enable dynamic cache read prefetch**

   CAUTION: Possible loss of data—Selecting the Enable write caching without batteries option lets write caching continue even when the batteries are discharged completely or are not fully charged. Typically, write caching is turned off temporarily by the RAID controller module until the batteries are charged. If you select this option and do not have a universal power supply for protection, you could lose data. In addition, you could lose data if you do not have RAID controller module batteries and you select the Enable write caching without batteries option.

   NOTE: When the Optional RAID controller module batteries option is enabled, the Enable write caching does not appear. The Enable write caching without batteries is still available, but it is not checked by default.

   NOTE: Cache is automatically flushed after the Enable write caching check box is disabled.
5. Click **OK**.

   A message prompts you to confirm the change in the virtual disk modification priority.
6. Click **Yes**.
7. Click **OK**.

   The **Change Virtual Disk Propreties - Progress** dialog is displayed.

Changing The Segment Size Of A Virtual Disk

You can change the segment size on a selected virtual disk. During this operation, I/O performance is affected, but your data remains available.

Follow these guidelines to proceed with changing the segment size:

- You cannot cancel this operation after it starts.
- Do not start this operation unless the disk group is in Optimal status.
- The MD Storage Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the menu. Allowed transitions usually are double or half of current segment size. For example, if the current virtual disk segment size is 32 KB, a new virtual disk segment size of either 16 KB or 64 KB is allowed.

NOTE: The operation to change the segment size is slower than other modification operations (for example, changing RAID levels or adding free capacity to a disk group). This slowness is the result of how the data is reorganized and the temporary internal backup procedures that occur during the operation.

The amount of time that a change segment size operation takes depends on:

- The I/O load from the host
- The modification priority of the virtual disk
- The number of physical disks in the disk group
- The number of physical disk ports
- The processing power of the storage array RAID controller modules

If you want this operation to complete faster, you can change the modification priority to the highest level, although this may decrease system I/O performance.

To change the segment size of a virtual disk:

1. In the AMW, select the **Storage & Copy Services** tab and select a virtual disk.
2. From the menu bar, select **Storage → Virtual Disk → Change → Segment Size.**
3. Select the required segment size.

   A message prompts you to confirm the selected segment size.
4. Click **Yes**.

   The segment size modification operation begins. The virtual disk icon in the Details pane shows an Operation in Progress status while the operation is taking place.

   NOTE: To view the progress or change the priority of the modification operation, select a virtual disk in the disk group, and from the menu bar, select **Storage → Virtual Disk → Change → Modification Priority.**

Changing The IO Type

You can specify the virtual disk I/O characteristics for the virtual disks that you are defining as part of the storage array configuration. The expected I/O characteristics of the virtual disk is used by the system to indicate an applicable default virtual disk segment size and dynamic cache read prefetch setting. See the online help topics for information on the Automatic Configuration Wizard.

NOTE: The dynamic cache read prefetch setting can be changed later by selecting **Storage → Virtual Disk → Change → Cache Settings** from the menu bar. You can change the segment size later by selecting **Storage → Virtual Disk → Change → Segment Size** from the menu bar.

The I/O characteristic types shown below are only presented during the create virtual disk process.

When you choose one of the virtual disk I/O characteristics, the corresponding dynamic cache prefetch setting and segment size that are typically well suited for expected I/O patterns are populated in the **Dynamic cache read prefetch** field and the **Segment size** field.

To change the I/O type:

1. To enable read caching, select **Enable read caching**.
2. To enable dynamic cache read prefetch, select **Enable dynamic cache read prefetch**.
3. To enable write caching, select **Enable write caching**.
4. Select one of the following:
   - **Enable write caching with mirroring** — Select this option to mirror cached data across two redundant RAID controller modules that have the same cache size.
   - **Enable write caching without batteries** — Select this option to permit write caching to continue even if the RAID controller module batteries are discharged completely, not fully charged, or are not present.

   NOTE: Cache is automatically flushed if you disable **Enable write caching**.
5. Click **OK**.
6. In the confirmation dialog, click **Yes**.

   A progress dialog is displayed, which indicates the number of virtual disks being changed.

Thin Virtual Disks

When creating virtual disks from a disk pool, you have the option to create thin virtual disks instead of standard virtual disks. Thin virtual disks are created with physical (or preferred) and virtual capacity, allowing flexibility to meet increasing capacity requirements.

When you create standard virtual disks, you allocate all available storage based on an estimation of how much space you need for application data and performance. If you want to expand the size of a standard virtual disk in the future, you must add physical disks to your existing disk groups or disk pools. Thin volumes allow you to create large virtual disks with smaller physical storage allocations that can be increased as required.

NOTE: Thin virtual disks can only be created from an existing disk pool.

Advantages Of Thin Virtual Disks

Thin virtual disks, also known as thin provisioning, present a more logical storage view to hosts.

Thin virtual disks allow you to dynamically allocate storage to each virtual disk as data is written. Using thin provisioning helps to eliminate large amounts of unused physical capacity that often occurs when creating standard virtual disks.

However, in certain cases, standard virtual disks may provide a more suitable alternative compared to thin provisioning, such as in situations when:

- you anticipate that storage consumption on a virtual disk is highly unpredictable or volatile
- an application relying on a specific virtual disk is exceptionally mission critical

Physical Vs Virtual Capacity On A Thin Virtual Disk

When you configure a thin virtual disk, you can specify the following types of capacity:

- physical (or preferred)
- virtual

Virtual capacity is capacity that is reported to the host, while physical capacity is the amount of actual physical disk space allocated for data write operations. Generally, physical capacity is much smaller than virtual capacity.

Thin provisioning allows virtual disks to be created with a large virtual capacity but a relatively small physical capacity. This is beneficial for storage utilization and efficiency because it allows you to increase capacity as application needs change, without disrupting data throughput. You can also set a utilization warning threshold that causes MD Storage Manager to generate an alert when a specified percentage of physical capacity is reached.

Changing Capacity On Existing Thin Virtual Disks

If the amount of space used by the host for read/write operations (sometimes called consumed capacity) exceeds the amount of physical capacity allocated on a standard virtual disk, the storage array cannot accommodate additional write requests until the physical capacity is increased. However, on a thin virtual disk, MD Storage Manager can automatically expand physical capacity of a thin virtual disk. You can also do it manually using **Storage → Virtual Disk → Increase Repository Capacity.** If you select the automatic expansion option, you can also set a maximum expansion capacity. The maximum expansion capacity enables you to limit the automatic growth of a virtual disk to an amount less than the defined virtual capacity.

NOTE: Since less than full capacity is allocated when you create a thin virtual disk, insufficient free capacity may exist when certain operations are performed, such as snapshot images and snapshot virtual disks. If this occurs, an alert threshold warning is displayed.

Thin Virtual Disk Requirements And Limitations

The following table provides the minimum and maximum capacity requirements applicable to thin virtual disks.

Table 3. Minimum and Maximum Capacity Requirements

Capacity Types	Size	
Virtual capacity		
	Minimum	32 MB
	Maximum	63 TB
Physical capacity		
	Minimum	4 GB
	Maximum	64 TB

The following limitations apply to thin virtual disks:

- The segment size of a thin virtual disk cannot be changed.
- The pre-read consistency check for a thin virtual disk cannot be enabled.
- A thin virtual disk cannot serve as the target virtual disk in a Virtual Disk Copy.
- A thin virtual disk cannot be used in a Snapshot (Legacy) operation.
- A thin virtual disk cannot be used in a Remote Replication (Legacy) operation.

Thin Volume Attributes

When you create a thin virtual disk from free capacity in an existing disk pool, you can manually set disk attributes or allow MD Storage Manager to assign default attributes. The following manual attributes are available:

- **Preferred Capacity** — Sets the initial physical capacity of the virtual disk (MB, GB or TB). Preferred capacity in a disk pool is allocated in 4 GB increments. If you specify a capacity amount that is not a multiple of 4 GB, MD

Storage Manager assigns a 4 GB multiple and assigns the remainder as unused. If space exists that is not a 4 GB multiple, you can use it to increase the size of the thin virtual disk. To increase the size of the thin virtual disk, select **Storage → Virtual Disk → Increase Capacity.**

- **Repository Expansion Policy** — Select either **Automatic** or **Manual** to indicate whether MD Storage Manager must automatically expand physical capacity thresholds. If you select **Automatic**, enter a **Maximum Expansion Capacity** value that triggers automatic capacity expansion. The MD Storage Manager expands the preferred capacity in increments of 4 GB until it reaches the specified capacity. If you select **Manual**, automatic expansion does not occur and an alert is displayed when the **Warning Threshold** value percentage is reached.
- **Warning Threshold** — When consumed capacity reaches the specified percentage, MD Storage Manager sends an E-mail or SNMP alert.

Thin Virtual Disk States

The following are the virtual disk states displayed in MD Storage Manager:

- **Optimal** — Virtual disk is operating normally.
- **Full** — Physical capacity of a thin virtual disk is full and no more host write requests can be processed.
- **Over Threshold** — Physical capacity of a thin virtual disk is at or beyond the specified **Warning Threshold** percentage. The storage array status is shown as **Needs Attention**.
- **Failed** — Virtual disk failed, and is no longer available for read or write operations. The storage array status is shown as **Needs Attention**.

Comparison—Types Of Virtual Disks And Copy Services

The availability of copy services depends on the type of virtual disk that you are working with.

The following table shows the copy services features supported on each type of virtual disk.

Copy Services Feature	Standard Virtual Disk in a Disk Group	Standard Virtual Disk in a Disk Pool	Thin Virtual Disk
Snapshot (Legacy)	Supported	Not supported	Not supported
Snapshot image	Supported	Supported	Supported
Snapshot virtual disk	Supported	Supported	Supported
Rollback of snapshot	Supported	Supported	Supported
Delete virtual disk with snapshot images or snapshot virtual disks	Supported	Supported	Supported
Consistency group membership	Supported	Supported	Supported
Remote Replication (Legacy)	Supported	Not supported	Not supported
Remote Replication	Supported	Supported	Not supported

The source of a virtual disk copy can be either a standard virtual disk in a disk group, a standard virtual disk in a disk pool, or a thin virtual disk. The target of a virtual disk copy can be only a standard virtual disk in a disk group or a standard virtual disk in a disk pool, not a thin virtual disk. The following table summarizes the types of virtual disks you can use in a virtual disk copy.

Virtual Disk Copy Source	Virtual Disk Copy Target	Availability
Standard virtual disk	Standard virtual disk	Supported
Thin virtual disk	Standard virtual disk	Supported
Standard virtual disk	Thin virtual disk	Not supported
Thin virtual disk	Thin virtual disk	Not supported

Rollback On Thin Virtual Disks

Rollback operations are fully supported on thin virtual disks. A rollback operation restores the logical content of a thin virtual disk to match the selected snapshot image. There is no change to the consumed capacity of the thin virtual disk as a result of a rollback operation.

Initializing A Thin Virtual Disk

CAUTION: Possible loss of data – Initializing a thin virtual disk erases all data from the virtual disk. If you have questions, contact your Technical Support representative before performing this procedure.

When a thin virtual disk is created, it is automatically initialized. However, the MD Storage Manger Recovery Guru may advise that you manually initialize a thin virtual disk to recover from certain failure conditions. If you choose to reinitialize a thin virtual disk, you have several options:

- Keep the same physical capacity — If you keep the same physical capacity, the virtual disk can keep its current repository virtual disk, which saves initialization time.
- Change the physical capacity — If you change the physical capacity, a new repository virtual disk is created and you can optionally change the repository expansion policy and warning threshold.
- Move the repository to a different disk pool.

Initializing a thin virtual disk erases all data from the virtual disk. However, host mappings, virtual capacity, repository expansion policy and security settings are preserved. Initialization also clears the block indices, which causes unwritten blocks to be read as if they are zero-filled. After initialization, the thin virtual disk appears to be completely empty.

The following types of virtual disks cannot be initialized:

- Snapshot (Legacy) virtual disk
- Base virtual disk of a Snapshot virtual disk
- Primary virtual disk in a Remote Replication relationship
- Secondary virtual disk in a Remote Replication relationship
- Source virtual disk in a Virtual Disk Copy
- Target virtual disk in a Virtual Disk Copy
- Thin virtual disk that already has an initialization in progress
- Thin virtual disk that is not in the **Optimal** state

Initializing A Thin Virtual Disk With The Same Physical Capacity

CAUTION: Initializing a thin virtual disk erases all data from the virtual disk.

- You can create thin virtual disks only from disk pools, not from disk groups.
- By initializing a thin virtual disk with the same physical capacity, the original repository is maintained but the contents of the thin virtual disk are deleted.

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the thin virtual disk that you want to initialize.

   The thin virtual disks are listed under the **Disk Pools** node.
3. Select **Storage → Virtual Disk → Advanced → Initialize.**

   The **Initialize Thin Virtual Disk** window is displayed.
4. Select **Keep existing repository**, and click **Finish**.

   The **Confirm Initialization of Thin Virtual Disk** window is displayed.
5. Read the warning and confirm if you want to initialize the thin virtual disk.
6. Type `yes`, and click **OK**.

   The thin virtual disk initializes.

Initializing A Thin Virtual Disk With A Different Physical Capacity

CAUTION: Initializing a thin virtual disk erases all data from the virtual disk.

- You can create thin virtual disks only from disk pools, not from disk groups.
- By initializing a thin virtual disk with the same physical capacity, the original repository is maintained but the contents of the thin virtual disk are deleted.

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the thin virtual disk that you want to initialize.

   The thin virtual disks are listed under the **Disk Pools** node.
3. Select **Storage → Virtual Disk → Advanced → Initialize.**

   The **Initialize Thin Virtual Disk** window is displayed.
4. Select **Use a different repository**.
5. Based on whether you want to keep the current repository for future use, select or clear **Delete existing repository**, and click **Next**.
6. Select one of the following:

   – Yes — If there more than one disk pool on your storage array

   – No — If there is only one disk pool on your storage array

   The **Select Disk Pool** window is displayed.
7. Select **Keep existing disk pool**, and click **Next**.

   The **Select Repository** window is displayed.
8. Use the **Preferred capacity** box to indicate the initial physical capacity of the virtual disk and the **Units** list to indicate the specific capacity units to use (MB, GB, or TB).

   NOTE: Do not allocate all of the capacity to standard virtual disks — ensure that you keep storage capacity for copy services (snapshots (legacy), snapshot images, snapshot virtual disks, virtual disk copies, and remote replications).

   NOTE: Regardless of the capacity specified, capacity in a disk pool is allocated in 4 GB increments. Any capacity that is not a multiple of 4 GB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4 GB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the virtual disk.

   Based on the value that you entered in the previous step, the **Disk pool physical capacity candidates** table is populated with matching repositories.
9. Select a repository from the table.

   Existing repositories are placed at the top of the list.

NOTE: The benefit of reusing an existing repository is that you can avoid the initialization process that occurs when you create a new one.

10. If you want to change the repository expansion policy or warning threshold, click **View advanced repository settings**.
    - **Repository expansion policy** – Select either **Automatic** or **Manual**. When the consumed capacity gets close to the physical capacity, you can expand the physical capacity. The MD storage management software can automatically expand the physical capacity or you can do it manually. If you select **Automatic**, you also can set a maximum expansion capacity. The maximum expansion capacity allows you to limit the virtual disk's automatic growth below the virtual capacity. The value for the maximum expansion capacity must be a multiple of 4 GB.
    - **Warning threshold** – In the **Send alert when repository capacity reaches** field, enter a percentage. The MD Storage Manager sends an alert notification when the physical capacity reaches the full percentage.
11. Click **Finish**.

    The **Confirm Initialization of Thin Virtual Disk** window is displayed.
12. Read the warning and confirm if you want to initialize the thin virtual disk.
13. Type `yes`, and click **OK**.

    The thin virtual disk initializes.

Initializing A Thin Virtual Disk And Moving It To A Different Disk Pool

CAUTION: Initializing a thin virtual disk erases all data from the virtual disk.

NOTE: You can create thin virtual disks only from disk pools, not from disk groups.

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the thin virtual disk that you want to initialize.

   The thin virtual disks are listed under the **Disk Pools** node.
3. Select **Storage → Virtual Disk → Advanced → Initialize.**

   The **Initialize Thin Virtual Disk** window is displayed.
4. Based on whether you want to keep the current repository for future use, select or clear **Delete existing repository**, and click **Next**.

   The **Select Disk Pool** window is displayed.
5. Select the **Select a new disk pool** radio button.
6. Select a new disk pool from the table, and click **Next**.

   The **Select Repository** window is displayed.
7. Select **Keep existing disk pool**, and click **Next**.

   The **Select Repository** window is displayed.
8. Use the **Preferred capacity** box to indicate the initial physical capacity of the virtual disk and the **Units** list to indicate the specific capacity units to use (MB, GB, or TB).

   NOTE: Do not allocate all of the capacity to standard virtual disks — ensure that you keep storage capacity for copy services (snapshots (legacy), snapshot images, snapshot virtual disks, virtual disk copies, and remote replications).

   NOTE: Regardless of the capacity specified, capacity in a disk pool is allocated in 4 GB increments. Any capacity that is not a multiple of 4 GB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4 GB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the virtual disk.

   Based on the value that you entered in the previous step, the **Disk pool physical capacity candidates** table is populated with matching repositories.

9. Select a repository from the table.

   Existing repositories are placed at the top of the list.

   NOTE: The benefit of reusing an existing repository is that you can avoid the initialization process that occurs when you create a new one.

10. If you want to change the repository expansion policy or warning threshold, click **View advanced repository settings**.

    - **Repository expansion policy** – Select either **Automatic** or **Manual**. When the consumed capacity gets close to the physical capacity, you can expand the physical capacity. The MD Storage Manager can automatically expand the physical capacity or you can do it manually. If you select **Automatic**, you also can set a maximum expansion capacity. The maximum expansion capacity allows you to limit the virtual disk's automatic growth below the virtual capacity. The value for the maximum expansion capacity must be a multiple of 4 GB.
    - **Warning threshold** – In the **Send alert when repository capacity reaches** field, enter a percentage. The MD Storage Manager sends an alert notification when the physical capacity reaches the full percentage.

11. Click **Finish**.

    The **Confirm Initialization of Thin Virtual Disk** window is displayed.

12. Read the warning and confirm if you want to initialize the thin virtual disk.

13. Type `yes`, and click **OK**.

    The thin virtual disk initializes.

Changing A Thin Virtual Disk To A Standard Virtual Disk

If you want to change a thin virtual disk to a standard virtual disk, use the Virtual Disk Copy operation to create a copy of the thin virtual disk. The target of a virtual disk copy must always be a standard virtual disk.

Choosing An Appropriate Physical Disk Type

You can create disk groups and virtual disks in the storage array. You must select the capacity that you want to allocate for the virtual disk from either unconfigured capacity, free capacity, or an existing disk pool available in the storage array. Then you define basic and optional advanced parameters for the virtual disk.

With the advent of different physical disk technologies, it is now possible to mix physical disks with different media types and different interface types within a single storage array.

Physical Disk Security With Self Encrypting Disk

Self Encrypting Disk (SED) technology prevents unauthorized access to the data on a physical disk that is physically removed from the storage array. The storage array has a security key. Self encrypting disks provide access to data only through an array that has the correct security key.

The self encrypting disk or a security capable physical disk encrypts data during writes and decrypts data during reads. For more information, see the online help topics.

You can create a secure disk group from security capable physical disks. When you create a secure disk group from security capable physical disks, the physical disks in that disk group become security enabled. When a security capable physical disk has been security enabled, the physical disk requires the correct security key from a RAID controller module to read or write the data. All of the physical disks and RAID controller modules in a storage array share the same security key. The shared security key provides read and write access to the physical disks, while the physical disk encryption key on each physical disk is used to encrypt the data. A security capable physical disk works like any other physical disk until it is security enabled.

Whenever the power is turned off and turned on again, all of the security enabled physical disks change to a security locked state. In this state, the data is inaccessible until the correct security key is provided by a RAID controller module.

You can view the self encrypting disk status of any physical disk in the storage array from the Physical Disk Properties dialog. The status information reports whether the physical disk is:

- Security capable
- Secure — Security enabled or disabled
- Read/Write Accessible — Security locked or unlocked

You can view the self encrypting disk status of any disk group in the storage array. The status information reports whether the storage array is:

- Security capable
- Secure

The following table shows how to interpret the security status of a disk group.

Table 4. Interpretation of Security Status of Disk Group

Secure	Security Capable - Yes	Security Capable - No
Yes	The disk group is composed of all SED physical disks and is in a Secure state.	Not applicable. Only SED physical disks can be in a Secure state.
No	The disk group is composed of all SED physical disks and is in a Non-Secure state.	The disk group is not entirely composed of SED physical disks.

The **Physical Disk Security** menu is displayed in the **Storage Array** menu. The **Physical Disk Security** menu has the following options:

- **Create Key**
- **Change Key**
- **Save Key**
- **Validate Key**
- **Import Key**
- **Unlock Drives**

NOTE: If you have not created a security key for the storage array, the **Create Key** option is active. If you have created a security key for the storage array, the **Create Key** option is inactive with a check mark to the left. The **Change Key** option, the **Save Key** option, and the **Validate Key** option are now active.

The **Secure Physical Disks** option is displayed in the **Disk Group** menu. The **Secure Physical Disks** option is active if these conditions are true:

- The selected storage array is not security enabled but is comprised entirely of security capable physical disks.
- The storage array contains no snapshot base virtual disks or snapshot repository virtual disks.
- The disk group is in an Optimal state.
- A security key is set up for the storage array.

NOTE: The **Secure Physical Disks** option is inactive if these conditions are not true.

The **Secure Physical Disks** option is inactive with a check mark to the left if the disk group is already security enabled.

The **Create a secure disk group** option is displayed in the **Create Disk Group Wizard–Disk Group Name** and **Physical Disk Selection** dialog. The **Create a secure disk group** option is active only when these conditions are met:

- A security key is installed in the storage array.
- At least one security capable physical disk is installed in the storage array.
- All of the physical disks that you selected on the **Hardware** tab are security capable physical disks.

You can erase security enabled physical disks so that you can reuse the drives in another disk group or in another storage array. When you erase security enabled physical disks, ensure that the data cannot be read. When all of the physical disks that you have selected in the Physical Disk type pane are security enabled, and none of the selected physical disks is part of a disk group, the **Secure Erase** option is displayed in the **Hardware** menu.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from self encrypting disk, and should not be confused with the pass phrase that is used to protect copies of a security key. However, it is good practice to set a storage array password.

Creating A Security Key

When you create a security key, it is generated by and securely stored by the array. You cannot read or view the security key. A copy of the security key must be kept on some other storage medium for backup in case of system failure or for transfer to another storage array. A pass phrase that you provide is used to encrypt and decrypt the security key for storage on other media.

When you create a security key, you also provide information to create a security key identifier. Unlike the security key, you can read or view the security key identifier. The security key identifier is also stored on a physical disk or transportable media. The security key identifier is used to identify which key the storage array is using.

To create a security key:

1. In the AMW, from the menu bar, select **Storage Array** → **Security** → **Physical Disk Security** → **Create Key** .
2. Perform one of these actions:
   - If the **Create Security Key** dialog is displayed, go to step 6.
   - If the **Storage Array Password Not Set** or **Storage Array Password Too Weak** dialog is displayed, go to step 3.
3. Choose whether to set (or change) the storage array password at this time.
   - Click **Yes** to set or change the storage array password. The **Change Password** dialog is displayed. Go to step 4.
   - Click **No** to continue without setting or changing the storage array password. The **Create Security Key** dialog is displayed. Go to step 6.
4. **In New password**, enter a string for the storage array password. If you are creating the storage array password for the first time, leave **Current password** blank. Follow these guidelines for cryptographic strength when you create the storage array password:
   - The password should be between eight and 30 characters long.
   - The password should contain at least one uppercase letter.
   - The password should contain at least one lowercase letter.
   - The password should contain at least one number.
   - The password should contain at least one non-alphanumeric character, for example, < > @ +.
5. In **Confirm new password**, re-enter the exact string that you entered in **New password**.
6. **In Security key identifier**, enter a string that becomes part of the secure key identifier.

   You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols. Additional characters are generated automatically and is appended to the end of the string that you enter. The generated characters help to ensure that the secure key identifier is unique.
7. Enter a path and file name to save the security key file by doing one of the following:

- Edit the default path by adding a file name to the end of the path.
- Click **Browse** to navigate to the required folder, then add a file name to the end of the path.

8. In **Pass phrase** dialog box, enter a string for the pass phrase.

   The pass phrase must:

   - be between eight and 32 characters long
   - contain at least one uppercase letter
   - contain at least one lowercase letter
   - contain at least one number
   - contain at least one non-alphanumeric character, for example, < > @ +

   The pass phrase that you enter is masked.

   NOTE: Create Key is active only if the pass phrase meets the above mentioned criterion.

9. In the **Confirm pass phrase** dialog box, re-enter the exact string that you entered in the **Pass phrase** dialog box.

   Make a record of the pass phrase that you entered and the security key identifier that is associated with the pass phrase. You need this information for later secure operations.

10. Click **Create Key**.

11. If the **Invalid Text Entry** dialog is displayed, select:

    - **Yes** — There are errors in the strings that were entered. The **Invalid Text Entry** dialog is displayed. Read the error message in the dialog, and click **OK**. Go to step 6.
    - **No** — There are no errors in the strings that were entered. Go to step 12.

12. Make a record of the security key identifier and the file name from the **Create Security Key Complete** dialog, and click **OK**.

After you have created a security key, you can create secure disk groups from security capable physical disks. Creating a secure disk group makes the physical disks in the disk group security enabled. Security enabled physical disks enter Security Locked status whenever power is re-applied. They can be unlocked only by a RAID controller module that supplies the correct key during physical disk initialization. Otherwise, the physical disks remain locked, and the data is inaccessible. The Security Locked status prevents any unauthorized person from accessing data on a security enabled physical disk by physically removing the physical disk and installing the physical disk in another computer or storage array.

Changing A Security Key

When you change a security key, a new security key is generated by the system. The new key replaces the previous key. You cannot view or read the key. However, a copy of the security key must be kept on some other storage medium for backup in case of system failure or for transfer to another storage array. A pass phrase that you provide encrypts and decrypts the security key for storage on other media. When you change a security key, you also provide information to create a security key identifier. Changing the security key does not destroy any data. You can change the security key at any time.

Before you change the security key, ensure that:

- All virtual disks in the storage array are in Optimal status.
- In storage arrays with two RAID controller modules, both are present and working normally.

To change the security key:

1. In the AMW menu bar, select **Storage Array** → **Security** → **Physical Disk Security** → **Change Key**.

The **Confirm Change Security Key** window is displayed.

2. Type `yes` in the text field, and click **OK**.

   The **Change Security Key** window is displayed.

3. In **Secure key identifier**, enter a string that become part of the secure key identifier.

   You may leave the text box blank, or enter up to 189 alphanumeric characters without white space, punctuation, or symbols. Additional characters is generated automatically.

4. Edit the default path by adding a file name to the end of the path or click **Browse**, navigate to the required folder, and enter the name of the file.

5. In **Pass phrase**, enter a string for the pass phrase.

   The pass phrase must meet the following criteria:

   - It must be between eight and 32 characters long.
   - It must contain at least one uppercase letter.
   - It must contain at least one lowercase letter.
   - It must contain at least one number.
   - It must contain at least one non-alphanumeric character (for example, < > @ +).

   The pass phrase that you enter is masked.

6. In **Confirm pass phrase**, re-enter the exact string you entered in **Pass phrase**.

   Make a record of the pass phrase you entered and the security key identifier it is associated with. You need this information for later secure operations.

7. Click **Change Key**.

8. Make a record of the security key identifier and the file name from the **Change Security Key Complete** dialog, and click **OK**.

Saving A Security Key

You save an externally storable copy of the security key when the security key is first created and each time it is changed. You can create additional storable copies at any time. To save a new copy of the security key, you must provide a pass phrase. The pass phrase you choose does not need to match the pass phrase used when the security key was created or last changed. The pass phrase is applied to the particular copy of the security key you are saving.

To save the security key for the storage array,

1. In the AMW toolbar, select **Storage Array → Security → Physical Disk Security → Save Key.**

   The **Save Security Key File - Enter Pass Phrase** window is displayed.

2. Edit the default path by adding a file name to the end of the path or click **Browse**, navigate to the required folder and enter the name of the file.

3. In **Pass phrase**, enter a string for the pass phrase.

   The pass phrase must meet the following criteria:

   - It must be between eight and 32 characters long.
   - It must contain at least one uppercase letter.
   - It must contain at least one lowercase letter.
   - It must contain at least one number.
   - It must contain at least one non-alphanumeric character (for example, < > @ +).

   The pass phrase that you enter is masked.

4. In **Confirm pass phrase**, re-enter the exact string you entered in **Pass phrase**.

Make a record of the pass phrase you entered. You need it for later secure operations.

5. Click **Save**.
6. Make a record of the security key identifier and the file name from the **Save Security Key Complete** dialog, and click **OK**.

Validate Security Key

A file in which a security key is stored is validated through the Validate Security Key dialog. To transfer, archive, or back up the security key, the RAID controller module firmware encrypts (or wraps) the security key and stores it in a file. You must provide a pass phrase and identify the corresponding file to decrypt the file and recover the security key.

Data can be read from a security enabled physical disk only if a RAID controller module in the storage array provides the correct security key. If security enabled physical disks are moved from one storage array to another, the appropriate security key must also be imported to the new storage array. Otherwise, the data on the security enabled physical disks that were moved is inaccessible.

See the online help topics for more information on validating the security key.

Unlocking Secure Physical Disks

You can export a security enabled disk group to move the associated physical disks to a different storage array. After you install those physical disks in the new storage array, you must unlock the physical disks before data can be read from or written to the physical disks. To unlock the physical disks, you must supply the security key from the original storage array. The security key on the new storage array is different and cannot unlock the physical disks.

You must supply the security key from a security key file that was saved on the original storage array. You must provide the pass phrase that was used to encrypt the security key file to extract the security key from this file.

For more information, see the online help topics.

Erasing Secure Physical Disks

In the AMW, when you select a security enabled physical disk that is not part of a disk group, the **Secure Erase** menu item is enabled on the Physical Disk menu. You can use the secure erase procedure to re-provision a physical disk. You can use the Secure Erase option if you want to remove all of the data on the physical disk and reset the physical disk security attributes.

CAUTION: Possible loss of data access—The Secure Erase option removes all of the data that is currently on the physical disk. This action cannot be undone.

Before you complete this option, make sure that the physical disk that you have selected is the correct physical disk. You cannot recover any of the data that is currently on the physical disk.

After you complete the secure erase procedure, the physical disk is available for use in another disk group or in another storage array. See the online help topics for more information on the secure erase procedure.

Configuring Hot Spare Physical Disks

Guidelines to configure host spare physical disks:

CAUTION: If a hot spare physical disk does not have Optimal status, follow the Recovery Guru procedures to correct the problem before you try to unassign the physical disk. You cannot assign a hot spare physical disk if it is in use (taking over for a failed physical disk).

- You can use only unassigned physical disks with **Optimal** status as hot spare physical disks.

- You can unassign only hot spare physical disks with **Optimal**, or **Standby** status. You cannot unassign a hot spare physical disk that has the **In Use** status. A hot spare physical disk has the **In Use** status when it is in the process of taking over for a failed physical disk.
- Hot spare physical disks must be of the same media type and interface type as the physical disks that they are protecting.
- If there are secure disk groups and security capable disk groups in the storage array, the hot spare physical disk must match the security capability of the disk group.
- Hot spare physical disks must have capacities equal to or larger than the used capacity on the physical disks that they are protecting.
- The availability of enclosure loss protection for a disk group depends on the location of the physical disks that comprise the disk group. To make sure that enclosure loss protection is not affected, you must replace a failed physical disk to initiate the copyback process. See Enclosure Loss Protection.

To assign or unassign hot spare physical disks:

1. In the AMW, select the **Hardware** tab.
2. Select one or more unassigned physical disks.
3. Perform one of these actions:
   - From the menu bar, select **Hardware** → **Hot Spare Coverage.**
   - Right-click the physical disk and select **Hot Spare Coverage** from the pop-up menu.

   The **Hot Spare Physical Disk Options** window is displayed.
4. Select the appropriate option, you can select:
   - **View/change current hot spare coverage** — to review hot spare coverage and to assign or unassign hot spare physical disks, if necessary. See step 5.
   - **Automatically assign physical disks** — to create hot spare physical disks automatically for the best hot spare coverage using available physical disks.
   - **Manually assign individual physical disks** — to create hot spare physical disks out of the selected physical disks on the **Hardware** tab.
   - **Manually unassign individual physical disks** — to unassign the selected hot spare physical disks on the **Hardware** tab. See step 12.

     NOTE: This option is available only if you select a hot spare physical disk that is already assigned.
5. To assign hot spares, in the **Hot Spare Coverage** window, select a disk group in the **Hot spare coverage** area.
6. Review the information about the hot spare coverage in the **Details** area.
7. Click **Assign**.

   The **Assign Hot Spare** window is displayed.
8. Select the relevant Physical disks in the **Unassigned physical disks** area, as hot spares for the selected disk and click **OK**.
9. To unassign hot spares, in the **Hot Spare Coverage** window, select physical disks in the **Hot spare physical disks** area.
10. Review the information about the hot spare coverage in the **Details** area.
11. Click **Unassign**.

    A message prompts you to confirm the operation.
12. Type `yes` and click **OK**.

Hot Spares And Rebuild

A valuable strategy to protect data is to assign available physical disks in the storage array as hot spares. A hot spare adds another level of fault tolerance to the storage array.

A hot spare is an idle, powered-on, stand-by physical disk ready for immediate use in case of disk failure. If a hot spare is defined in an enclosure in which a redundant virtual disk experiences a physical disk failure, a rebuild of the degraded virtual disk is automatically initiated by the RAID controller modules. If no hot spares are defined, the rebuild process is initiated by the RAID controller modules when a replacement physical disk is inserted into the storage array.

Global Hot Spares

The MD Series storage arrays support global hot spares. A global hot spare can replace a failed physical disk in any virtual disk with a redundant RAID level as long as the capacity of the hot spare is equal to or larger than the size of the configured capacity on the physical disk it replaces, including its metadata.

Hot Spare Operation

When a physical disk fails, the virtual disk automatically rebuilds using an available hot spare. When a replacement physical disk is installed, data from the hot spare is copied back to the replacement physical disk. This function is called copy back. By default, the RAID controller module automatically configures the number and type of hot spares based on the number and capacity of physical disks in your system.

A hot spare may have the following states:

- A standby hot spare is a physical disk that has been assigned as a hot spare and is available to take over for any failed physical disk.
- An in-use hot spare is a physical disk that has been assigned as a hot spare and is currently replacing a failed physical disk.

Hot Spare Drive Protection

You can use a hot spare physical disk for additional data protection from physical disk failures that occur in a RAID Level 1, or RAID Level 5 disk group. If the hot spare physical disk is available when a physical disk fails, the RAID controller module uses redundancy data to reconstruct the data from the failed physical disk to the hot spare physical disk. When you have physically replaced the failed physical disk, a copyback operation occurs from the hot spare physical disk to the replaced physical disk. If there are secure disk groups and security capable disk groups in the storage array, the hot spare physical disk must match the security capability of the disk group. For example, a non-security capable physical disk cannot be used as a hot spare for a secure disk group.

NOTE: For a security capable disk group, security capable hot spare physical disks are preferred. If security capable physical disks are not available, non-security capable physical disks may be used as hot spare physical disks. To ensure that the disk group is retained as security capable, the non-security capable hot spare physical disk must be replaced with a security capable physical disk.

If you select a security capable physical disk as hot spare for a non-secure disk group, a dialog box is displayed indicating that a security capable physical disk is being used as a hot spare for a non-secure disk group.

The availability of enclosure loss protection for a disk group depends on the location of the physical disks that comprise the disk group. The enclosure loss protection might be lost because of a failed physical disk and location of the hot spare physical disk. To make sure that enclosure loss protection is not affected, you must replace a failed physical disk to initiate the copyback process.

The virtual disk remains online and accessible while you are replacing the failed physical disk, because the hot spare physical disk is automatically substituted for the failed physical disk.

Enclosure Loss Protection

Enclosure loss protection is an attribute of a disk group. Enclosure loss protection guarantees accessibility to the data on the virtual disks in a disk group if a total loss of communication occurs with a single expansion enclosure. An example of total loss of communication may be loss of power to the expansion enclosure or failure of both RAID controller modules.

CAUTION: Enclosure loss protection is not guaranteed if a physical disk has already failed in the disk group. In this situation, losing access to an expansion enclosure and consequently another physical disk in the disk group causes a double physical disk failure and loss of data.

Enclosure loss protection is achieved when you create a disk group where all of the physical disks that comprise the disk group are located in different expansion enclosures. This distinction depends on the RAID level. If you choose to create a disk group by using the Automatic method, the software attempts to choose physical disks that provide enclosure loss protection. If you choose to create a disk group by using the Manual method, you must use the criteria specified below.

RAID Level	Criteria for Enclosure Loss Protection
RAID level 5 or RAID level 6	Ensure that all the physical disks in the disk group are located in different expansion enclosures. Because a RAID level 5 requires a minimum of three physical disks, enclosure loss protections cannot be achieved if your storage array has less than three expansion enclosures. Because a RAID level 6 requires a minimum of five physical disks, enclosure loss protections cannot be achieved if your storage array has less than five expansion enclosures.
RAID level 1	Ensure that each physical disk in a mirrored pair is located in a different expansion enclosure. This enables you to have more than two physical disks in the disk group within the same expansion enclosure. For example, if you are creating a six physical disk, disk group (three-mirrored pairs), you could achieve enclosure loss protection with only two expansion enclosures by specifying that the physical disk in each mirrored pair are located in separate expansion enclosures. This example shows this concept: • Mirror pair 1 — Physical disk in enclosure 1 slot 1 and physical disk in enclosure 2 slot 1. • Mirror pair 2 — Physical disk in enclosure 1 slot 2 and physical disk in enclosure 2 slot 2. • Mirror pair 3 — Physical disk in enclosure 1 slot 3 and physical disk in enclosure 2 slot 3. Because a RAID level 1 disk group requires a minimum of two physical disks, enclosure loss protections cannot be achieved if your storage array has less than two expansion enclosures.
RAID level 0	Because RAID level 0 does not have consistency, you cannot achieve enclosure loss protection.

Drawer Loss Protection

In expansion enclosures that contain drawer-based physical disks, a drawer failure can prevent access to data on the virtual disks of a disk group.

Drawer loss protection for a disk group is based on the location of the physical disks that comprise the disk group. In the event of a single drawer failure, data on the virtual disks in a disk group will remain accessible if drawer loss protection configuration is followed. In such as case, if a drawer fails and the disk group is protected, the disk group changes to Degraded status and the data remains accessible.

To configure your storage for drawer loss protection, make sure that physical disks that are part of a disk group are located in different drawers with respect to their RAID levels, as shown in the following table.

Table 5. Drawer Loss Protection Requirements for Different RAID Levels

RAID Level	Drawer Loss Protection Requirements
RAID Level 6	RAID Level 6 requires a minimum of five physical disks. Place all of the physical disks in different drawers or place a maximum of two physical disks in the same drawer and the remaining physical disks in different drawers.
RAID Level 5	RAID Level 5 requires a minimum of three physical disks. Place all of the physical disks in different drawers for a RAID Level 5 disk group. Drawer loss protection cannot be achieved for RAID Level 5 if more than one physical disk is placed in the same drawer.
RAID Level 1 and RAID Level 10	RAID Level 1 requires a minimum of two physical disks. Make sure that each physical disk in a remotely replicated pair is located in a different drawer. By locating each physical disk in a different drawer, you can have more than two physical disks of the disk group within the same drawer. For example, if you create a RAID Level 1 disk group with six physical disks (three replicated pairs), you can achieve the drawer loss protection for the disk group with only two drawers as shown in this example: Six-physical disk RAID Level 1 disk group: Replicated pair 1 = Physical disk located in enclosure 1, drawer 0, slot 0, and physical disk in enclosure 0, drawer 1, slot 0 Replicated pair 2 = Physical disk in enclosure 1, drawer 0, slot 1, and physical disk in enclosure 1, drawer 1, slot 1 Replicated pair 3 = Physical disk in enclosure 1, drawer 0, slot 2, and physical disk in enclosure 2, drawer 1, slot 2 RAID Level 10 requires a minimum of four physical disks. Make sure that each physical disk in a remotely replicated pair is located in a different drawer.
RAID Level 0	You cannot achieve drawer loss protection because the RAID Level 0 disk group does not have redundancy.

NOTE: If you create a disk group using the Automatic physical disk selection method, MD Storage Manager attempts to choose physical disks that provide drawer loss protection. If you create a disk group by using the Manual physical disk selection method, you must use the criteria that are specified in the previous table.

If a disk group already has a Degraded status due to a failed physical disk when a drawer fails, drawer loss protection does not protect the disk group. The data on the virtual disks becomes inaccessible.

Host-To-Virtual Disk Mapping

After you create virtual disks, you must map them to the host(s) connected to the array.

The following are the guidelines to configure host-to-virtual disk mapping:

- Each virtual disk in the storage array can be mapped to only one host or host group.
- Host-to-virtual disk mappings are shared between controllers in the storage array.
- A unique LUN must be used by a host group or host to access a virtual disk.
- Each host has its own LUN address space. MD Storage Manager permits the same LUN to be used by different hosts or host groups to access virtual disks in a storage array.
- All operating system do not have the same number of LUNs available.

- You can define the mappings on the **Host Mappings** tab in the AMW. See Using The Host Mappings Tab.

Creating Host-To-Virtual Disk Mappings

Guidelines to define the mappings:

- An access virtual disk mapping is not required for an out-of-band storage array. If your storage array is managed using an out-of-band connection, and an access virtual disk mapping is assigned to the Default Group, an access virtual disk mapping is assigned to every host created from the Default Group.
- Most hosts have 256 LUNs mapped per storage partition. The LUN numbering is from 0 through 255. If your operating system restricts LUNs to 127, and you try to map a virtual disk to a LUN that is greater than or equal to 127, the host cannot access it.
- An initial mapping of the host group or host must be created using the Storage Partitioning Wizard before defining additional mappings. See Storage Partitioning.

To create host to virtual disk mappings:

1. In the AMW, select the **Host Mappings** tab.
2. In the object tree, select:
   - Default Group
   - Undefined mappings node
   - Individual defined mapping
   - Host group
   - Host
3. From the menu bar, select **Host Mappings** → **LUN Mapping** → **Add.**
   The **Define Additional Mapping** window is displayed.
4. In **Host group or host**, select the appropriate host group or host.
   All defined hosts, host groups, and the default group are displayed in the list.

   NOTE: When configuring an iSCSI storage array, if a host or a host group is selected that does not have a SAS host bus adapter (SAS HBA) host port defined, a warning dialog is displayed.
5. In **Logical unit number**, select a LUN.
   The supported LUNs are 0 through 255.
6. Select the virtual disk to be mapped in the **Virtual Disk** area.
   The **Virtual Disk** area lists the names and capacity of the virtual disks that are available for mapping based on the selected host group or selected host.
7. Click **Add**.

   NOTE: The **Add** button is inactive until a host group or host, LUN, and virtual disk are selected.
8. To define additional mappings, repeat step 4 through step 7.

   NOTE: After a virtual disk has been mapped once, it is no longer available in the **Virtual Disk** area.
9. Click **Close**.
   The mappings are saved. The object tree and the **Defined Mappings** pane in the **Host Mappings** tab are updated to reflect the mappings.

Modifying And Removing Host-To-Virtual Disk Mapping

You can modify or remove a host-to-virtual disk mapping for several reasons, such as an incorrect mapping or reconfiguration of the storage array. Modifying or removing a host-to-virtual disk mapping applies to both hosts and host groups.

CAUTION: Before you modify or remove a host-to-virtual disk mapping, stop any data access (I/O) to the virtual disks to prevent data loss.

To modify or remove host to virtual disk mapping:

1. In the AMW, select the **Host Mappings** tab.
2. In the **Defined Mappings** pane, perform one of these actions:
   - Select a single virtual disk, and select **Host Mappings → LUN Mapping → Change.**
   - Right-click the virtual disk, and select **Change** from the pop-up menu.
3. In the **Host group or host** list, select the appropriate host group or host.

   By default, the drop-down list shows the current host group or the host associated with the selected virtual disk.
4. In **Logical unit number**, select the appropriate LUN.

   The drop down list shows only the currently available LUNs that are associated with the selected virtual disk.
5. Click **OK**.

   NOTE: Stop any host applications associated with this virtual disk, and unmount the virtual disk, if applicable, from your operating system.
6. In the **Change Mapping** dialog, click **Yes** to confirm the changes.

   The mapping is checked for validity and is saved. The **Defined Mappings** pane is updated to reflect the new mapping. The object tree is also updated to reflect any movement of host groups or hosts.
7. If a password is set on the storage array, the **Enter Password** dialog is displayed. Type the current password for the storage array, and click **OK**.
8. If configuring a Linux host, run the **rescan_dm_devs** utility on the host, and remount the virtual disk if required.

   NOTE: This utility is installed on the host as part of the MD Storage Manager install process.
9. Restart the host applications.

Changing Controller Ownership Of The Virtual Disk

If the host has a single data-path to the MD storage array, the virtual disk must be owned by the controller to which the host is connected. You must configure this storage array before you start I/O operations and after the virtual disk is created. You can change the RAID controller module ownership of a standard virtual disk or a snapshot repository virtual disk. You cannot directly change the RAID controller module ownership of a snapshot virtual disk because the snapshot virtual disk inherits the RAID controller module owner of its associated source virtual disk. Changing the RAID controller module ownership of a virtual disk changes the preferred RAID controller module ownership of the virtual disk.

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. Sometimes both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts. Therefore, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped, ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed. Under certain operating system environments, it may be necessary to reconfigure the multi-path driver before an I/O path can be used.

To change the ownership of the virtual disk to the connected controller:

1. In the AMW, select the **Storage & Copy Services** tab and select a virtual disk.
2. From the menu bar, select the appropriate RAID controller module slot in **Storage → Virtual Disk → Change → Ownership/Preferred Path.**
3. Click **Yes** to confirm the selection.

Removing Host-To-Virtual Disk Mapping

To remove the host to virtual disk mapping:

1. In the AMW, select the **Host Mappings** tab.
2. Select a virtual disk under **Defined Mappings**.
3. Perform one of these actions:
   - From the menu bar, select **Host Mappings → LUN Mapping → Remove.**
   - Right-click the virtual disk, and select **Remove** from the pop-up menu.
4. Click **Yes** to remove the mapping.

Changing The RAID Controller Module Ownership Of A Disk Group

You can change the RAID controller module ownership of a disk group. You can also change the RAID controller module ownership of a standard virtual disk or a snapshot repository virtual disk. You cannot directly change the RAID controller module ownership of a snapshot virtual disk because the snapshot virtual disk inherits the RAID controller module owner of its associated source virtual disk. Changing the RAID controller module ownership of a virtual disk changes the preferred RAID controller module ownership of the virtual disk.

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. Sometimes both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts. Therefore, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped, ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed. Under certain operating system environments, it may be necessary to reconfigure the multi-path driver before an I/O path can be used.

To change the RAID controller module ownership of a disk group:

1. In the AMW, select the **Storage & Copy Services** tab and select a disk group.
2. From the menu bar, select **Storage → Disk Group → Change → Ownership/Preferred Path.**
3. Select the appropriate RAID controller module slot and click **Yes** to confirm the selection.

   CAUTION: Possible loss of data access—Changing ownership at the disk group level causes every virtual disk in that disk group to transfer to the other RAID controller module and use the new I/O path. If you do not want to set every virtual disk to the new path, change ownership at the virtual disk level instead.

   The ownership of the disk group is changed. I/O to the disk group is now directed through this I/O path.

   NOTE: The disk group may not use the new I/O path until the multi-path driver reconfigures and recognizes the new path. This action usually takes less than 5 minutes.

Changing The RAID Level Of A Disk Group

Changing the RAID level of a disk group changes the RAID levels of every virtual disk that comprises the disk group. Performance may be slightly affected during the operation.

Guidelines to change the RAID level of a disk group:

- You cannot cancel this operation after it begins.
- The disk group must be in Optimal status before you can perform this operation.
- Your data remains available during this operation.
- If you do not have enough capacity in the disk group to convert to the new RAID level, an error message is displayed, and the operation does not continue. If you have unassigned physical disks, use the **Storage → Disk Group → Add Physical Disks (Capacity)** option to add additional capacity to the disk group and then retry the operation.

To change the RAID level of a disk group:

1. In the AMW, select the **Storage & Copy Services** tab and select a disk group.
2. From the menu bar, select **Storage → Disk Group → Change → RAID Level.**
3. Select the appropriate RAID level and click **Yes** to confirm the selection.
   The RAID level operation begins.

Removing A Host-To-Virtual Disk Mapping Using Linux DMMP

To remove a host-to-virtual disk mapping using Linux DMMP:

1. Unmount the file system containing the virtual disk.
   Using the following command: `# umount filesystemDirectory`
2. Run the following command to display multi-pathing topology:
   `# multipath -ll`

   NOTE: Use the `multipath -ll` command:

   - If a new LUN is mapped, the new LUN is detected and given a multipathing device node.
   - If you increased volume capacity, the new capacity is displayed.

   NOTE: The virtual disk that you want to delete from the mapping. For example, the following information may be displayed:

```
mpath6 (3600a0b80000fb6e50000000e487b02f5) dm-10
DELL, MD32xx
[size=1.6T][features=3 queue_if_no_path
pg_init_retries 50][hwhandler=1 rdac]
\_ round-robin 0 [prio=6][active]
\_ 1:0:0:2 sdf 8:80 [active][ready]
\_ round-robin 0 [prio=1][enabled]
\_ 0:0:0:2 sde 8:64 [active][ghost]
```

In this example, the mpath6 device contains two paths:

```
-- /dev/sdf at Host 1, Channel 0, Target 0, LUN 2
--/dev/sde at Host 0, Channel 0, Target 0, LUN 2
```

3. Flush the multi-pathing device mapping using the following command:

   ```
   # multipath -f /dev/mapper/mapth_x
   ```

   Where, `mapth_x` is the device you want to delete.

4. Delete the paths related with this device using the following command:

   ```
   # echo 1 > /sys/block/sd_x/device/delete
   ```

   Where, `sd_x` is the SD node (disk device) returned by the `multipath` command. Repeat this command for all paths related to this device. For example:

   ```
   #echo 1 > /sys/block/sdf/device/delete
   #echo 1 > /sys/block/sde/device/delete
   ```

5. Remove mapping from c, or delete the LUN if necessary.
6. If you want to map another LUN or increase volume capacity, perform this action from MD Storage Manager.

   NOTE: If you are only testing LUN removal, you can stop at this step.

7. If a new LUN is mapped or volume capacity is changed, run the following command: `# rescan_dm_devs`

Restricted Mappings

Many hosts are able to map up to 256 LUNs (0 to 255) per storage partition. However, the maximum number of mappings differs because of operating system variables, failover driver issues, and potential data problems. The hosts listed in the table have these mapping restrictions.

If you try to map a virtual disk to a LUN that exceeds the restriction on these operating systems, the host is unable to access the virtual disk.

Operating System	Highest LUN
Windows Server 2003 and Windows Server 2008	255
Linux	255

Guidelines when you work with host types with LUN mapping restrictions:

- You cannot change a host adapter port to a restricted host type if there are already mappings in the storage partition that would exceed the limit imposed by the restricted host type.
- Consider the case of the Default Group that has access to LUNs up to 256 (0 to 255) and a restricted host type is added to the Default Group. In this case, the host that is associated with the restricted host type is able to access virtual disks in the Default Group with LUNs within its limits. For example, if the Default Group had two virtual disks mapped to LUNs 254 and 255, the host with the restricted host type would not be able to access those two virtual disks.
- If the Default Group has a restricted host type assigned and the storage partitions are disabled, you can map only a total of 32 LUNs. Any additional virtual disks that are created are put in the Unidentified Mappings area. If additional mappings are defined for one of these Unidentified Mappings, the **Define Additional Mapping** dialog shows the LUN list, and the **Add** button is unavailable.
- Do not configure dual mappings on a Windows host.
- If there is a host with a restricted host type that is part of a specific storage partition, all of the hosts in that storage partition are limited to the maximum number of LUNs allowed by the restricted host type.
- You cannot move a host with a restricted host type into a storage partition that already has LUNs mapped that are greater than what is allowed by the restricted host type. For example, if you have a restricted host type that

allows only LUNs up to 31, you cannot move that restricted host type into a storage partition that has LUNs greater than 31 already mapped.

The Default Group on the **Host Mappings** tab has a default host type. To change the host type, right-click on the host and select **Change Default Host Operating** System from the pop-up menu. If you set the default host type to a host type that is restricted, the maximum number of LUNs allowed in the Default Group for any host is restricted to the limit imposed by the restricted host type. If a particular host with a non-restricted host type becomes part of a specific storage partition, you are able to change the mapping to a higher LUN.

Storage Partitioning

A storage partition is a logical entity consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group. The first time you map a virtual disk to a specific host or host group, a storage partition is created. Subsequent virtual disk mappings to that host or host group do not create another storage partition.

One storage partition is sufficient if:

- Only one attached host accesses all of the virtual disks in the storage array
- All attached hosts share access to all of the virtual disks in the storage array

When you choose this type of configuration, all of the hosts must have the same operating system and special software (such as clustering software) to manage virtual disk sharing and accessibility.

More than one storage partition is required if:

- Specific hosts must access specific virtual disks in the storage array
- Hosts with different operating systems are attached to the same storage array. In this case, a storage partition is created for each host type

You can use the Storage Partitioning Wizard to define a single storage partition. The Storage Partitioning Wizard guides you through the major steps required to specify which host groups, hosts, virtual disks, and associated logical unit numbers (LUNs) are to be included in the storage partition.

Storage partitioning fails when:

- All mappings are defined
- You create a mapping for a host group that conflicts with an established mapping for a host in the host group
- You create a mapping for a host in a host group that conflicts with an established mapping for the host group

Storage partitioning is unavailable when:

- No valid host groups or hosts exist in the object tree on the **Host Mappings** tab
- No host ports are defined for the host being included in the storage partition
- All mappings are defined

NOTE: You can include a secondary virtual disk in a storage partition. However, any hosts that are mapped to the secondary virtual disk has read-only access until the virtual disk is promoted to a primary virtual disk, or the mirror relationship is removed.

Storage partitioning topology is the collection of elements, such as Default Group, host groups, hosts, and host ports shown as nodes in the object tree of the **Host Mappings** tab in the AMW. For more information, see Using The Host Mappings Tab.

If a storage partitioning topology is not defined, an informational dialog is displayed each time you select the **Host Mappings** tab. You must define the storage partitioning topology before you define the actual storage partition.

Disk Group And Virtual Disk Expansion

Adding free capacity to a disk group is achieved by adding unconfigured capacity on the array to the disk group. Data is accessible on disk groups, virtual disks, and physical disks throughout the entire modification operation. The additional free capacity can then be used to perform a virtual disk expansion on a standard or snapshot repository virtual disk.

Disk Group Expansion

To add free capacity to a disk group:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select a disk group.
3. From the menu bar, select **Storage → Disk Group → Add Physical Disks (Capacity).** Alternatively, right-click on the disk group and select **Add Physical Disks (Capacity)** from the pop-up menu.

   The **Add Free Capacity** window is displayed. Based on the RAID level, and the enclosure loss protection of the current disk group, a list of unassigned physical disks is displayed.

   NOTE: If the RAID level of the disk group is RAID Level 5, or RAID Level 6, and the expansion enclosure has enclosure loss protection, **Display only physical disks that will ensure enclosure loss protection** is displayed and is selected by default.
4. In the **Available physical disks** area, select physical disks up to the allowed maximum number of physical disks.

   NOTE: You cannot mix different media types or different interface types within a single disk group or virtual disk.
5. Click **Add**.

   A message prompts you to confirm your selection.
6. To add the capacity to the disk group, click **Yes**.

   NOTE: You can also use the Command Line Interface (CLI) on both Windows and Linux hosts to add free capacity to a disk group. For more information, see the CLI Guide.

   NOTE: After the capacity expansion is completed, additional free capacity is available in the disk group for creation of new virtual disks or expansion of existing virtual disks.

Virtual Disk Expansion

Virtual disk expansion is a dynamic modification operation that increases the capacity of standard virtual disks.

NOTE: Snapshot repository virtual disks can be expanded from the CLI or from MD Storage Manager. All other virtual disk types are expandable only from the CLI.

If you receive a warning that the snapshot repository virtual disk is becoming full, you may expand the snapshot repository virtual disk from MD Storage Manager. For step-by-step instructions, see Snapshot Repository Capacity.

Using Free Capacity

You can increase the capacity of a virtual disk using the free capacity on the disk group of the standard virtual disk or the snapshot repository virtual disk.

The Total Unconfigured Capacity node, shown in the **Storage & Copy Services** tab, is a contiguous region of unassigned capacity on a defined disk group. When increasing virtual disk capacity, some or all of the free capacity may be used to

achieve the required final capacity. Data on the selected virtual disk remains accessible while the process for increasing virtual disk capacity is in progress.

Using Unconfigured Capacity

You can increase the capacity of a standard virtual disk or a snapshot repository virtual disk using the unconfigured capacity when no free capacity exists on a disk group. An increase is achieved by adding unconfigured capacity, in the form of unassigned physical disks, to the disk group of the standard virtual disk or the snapshot repository virtual disk. See Disk Group Expansion.

For more information, see the online help topics.

Disk Group Migration

Disk group migration allows to you to migrate a disk group by exporting a disk group and then importing it to another storage array. You can also export a disk group to store data offline.

When you export a disk group, all the physical disks become offline. To ensure that the export is successful, at least two physical disks that are not part of the disk group you are migrating must be present in the storage array.

When you migrate the exported disk group to the new storage array, the import fails if a majority of the physical disks are not present in the group. For example, both the physical disks in a two-disk RAID 1 configuration, or three physical disks (one from each disk pair) in a four-disk RAID 10 configuration must be present.

Export Disk Group

The export disk group operation prepares the physical disks in the disk group for removal. You can remove the physical disks for offline storage, or you can import the disk group to a different storage array. After you complete the export disk group operation, all of the physical disks are offline. Any associated virtual disks or free capacity nodes are no longer shown in the MD Storage Manager.

Non-Exportable Components

You must remove or clear any non-exportable settings before you can complete the export disk group procedure. Remove or clear the following settings:

- Persistent reservations
- Host-to-virtual disk mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks
- Remote mirror pairs
- Mirror repositories

Exporting A Disk Group

On the source storage array:

1. Save the storage array configuration.
2. Stop all I/O, and unmount or disconnect the file systems on the virtual disks in the disk group.
3. Back up the data on the virtual disks in the disk group.
4. Locate the disk group, and label the physical disks.
5. Place the disk group offline.
6. Obtain blank physical disk modules or new physical disks.

On the target storage array, verify that:

- The target storage array has available physical disk slots.
- The target storage array supports the physical disks that you import.
- The target storage array can support the new virtual disks.
- The latest version of firmware is installed on the RAID controller module.

Import Disk Group

The import disk group operation adds the imported disk group to the target storage array. After you complete the import disk group operation, all of the physical disks have Optimal status. Any associated virtual disks or free capacity nodes are now shown in the MD Storage Manager installed on the target storage array.

NOTE: You lose access to your data during the export/import process.

NOTE: You must export a disk group before you move the disk group or import the disk group.

Importing A Disk Group

NOTE: You must insert all of the physical disks that are part of the disk group into the enclosure before the disk group can be imported.

The following settings are removed/cleared during the procedure:

- Persistent reservations
- Host-to-virtual disk mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks
- Remote mirror pairs
- Mirror repositories

On the target storage array:

1. Insert the exported physical disks into the available physical disk slots.
2. Review the Import Report for an overview of the disk group that you are importing.
3. Check for non-importable components.
4. Confirm that you want to proceed with the import procedure.

   NOTE: Some settings cannot be imported during the import disk group procedure.

Non-Importable Components

Some components cannot be imported during the import disk group procedure. These components are removed during the procedure:

- Persistent reservations
- Mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks

Storage Array Media Scan

The media scan is a background operation that examines virtual disks to verify that data is accessible. The process finds media errors before normal read and write activity is disrupted and reports errors to the event log.

NOTE: You cannot enable background media scans on a virtual disk comprised of Solid State Disks (SSDs).

Errors discovered by the media scan include:

- Unrecovered media error — Data could not be read on the first attempt or on any subsequent attempts. For virtual disks with redundancy protection, data is reconstructed, rewritten to the physical disk, and verified and the error is reported to the event log. For virtual disks without redundancy protection (RAID 1, RAID 5, and RAID 6 virtual disks), the error is not corrected but is reported to the event log.
- Recovered media error — Data could not be read by the physical disk on the first attempt but was successfully read on a subsequent attempt. Data is rewritten to the physical disk and verified and the error is reported to the event log.
- Redundancy mismatches error — The first 10 redundancy mismatches that are found on the virtual disk are reported to the event log.
- Unfixable error — Data could not be read and parity or redundancy information could not be used to regenerate the data. For example, redundancy information cannot be used to reconstruct the data on a degraded virtual disk. The error is reported to the event log.

For more information, see the online help topics.

Changing Media Scan Settings

To change the media scan settings:

1. In the AMW, select the **Storage & Copy Services** tab and select any virtual disk.
2. From the menu bar, select **Storage → Virtual Disk → Change → Media Scan Settings.**

   The **Change Media Scan Settings** window is displayed.
3. Deselect **Suspend media scan**, if selected.
4. In **Scan duration (in days)**, enter or select the duration (in days) for the media scan.

   The media scan duration specifies the number of days for which the media scan runs on the selected virtual disks.
5. To disable media scans on an individual virtual disk, select the virtual disk in the **Select virtual disks to scan** area, and deselect **Scan selected virtual disks**.
6. To enable media scans on an individual virtual disk, select the virtual disk in the **Select virtual disks to scan** area, and select **Scan selected virtual disks**.
7. To enable or disable the consistency check, select either **With consistency check** or **Without consistency check**.

   NOTE: A consistency check scans the data blocks in a RAID Level 5 virtual disk, or a RAID Level 6 virtual disk and checks the consistency information for each block. A consistency check compares data blocks on RAID Level 1 mirrored physical disks. RAID Level 0 virtual disks have no data consistency.
8. Click **OK**.

Suspending The Media Scan

You cannot perform a media scan while performing another long-running operation on the disk drive such as reconstruction, copy-back, reconfiguration, virtual disk initialization, or immediate availability formatting. If you want to perform another long-running operation, you should suspend the media scan.

NOTE: A background media scan is the lowest priority of the long-running operations.

To suspend a media scan:

1. In the AMW, select the **Storage & Copy Services** tab and select any virtual disk.
2. From the menu bar, select **Storage → Virtual Disk → Change → Media Scan Settings.**
   The **Change Media Scan Settings** window is displayed.
3. Select **Suspend media scan**.

   NOTE: This applies to all the virtual disks on the disk group.
4. Click **OK**.

8

Disk Pools And Disk Pool Virtual Disks

Disk pooling allows you to distribute data from each virtual disk randomly across a set of physical disks. Disk pooling provides RAID protection and consistent performance across a set of physical disks logically grouped together in the storage array. Although there is no limit on the maximum number of physical disks that can comprise a disk pool, each disk pool must have a minimum of 11 physical disks. Additionally, the disk pool cannot contain more physical disks than the maximum limit for each storage array. The physical disks in each disk pool must:

- be SAS or nearline SAS
- have the same physical disk speed (RPM)

  NOTE: The maximum physical disk speed is 15,000 rpm for standard SAS and 7,500 rpm for 3.5" nearline SAS.

  NOTE: In a disk pool, the physical disks must have the same capacities. If the physical disks have different capacities, the MD Storage Manager uses the smallest capacity among the physical disks in the pool. For example, if your disk pool is comprised of several 4 GB physical disks and several 8 GB physical disks, only 4 GB on each physical disk is utilized.

The data and redundancy information in a disk pool is distributed across all of the physical disks in the pool and provides the following benefits:

- Simplified configuration
- Better utilization of physical disks
- Reduced maintenance
- the ability to use thin provisioning

Difference Between Disk Groups And Disk Pools

Similar to a disk group, you can create one or more virtual disks in a disk pool. However, the disk pool differs from a disk group in the way data is distributed across the physical disks comprising the pool.

In a disk group, data is distributed across the physical disks based on RAID level. You can specify a RAID level when you create the disk group, then the data for each virtual disk is written sequentially across the set of physical disks comprising the disk group.

NOTE: Because disk pools can co-exist with disk groups, a storage array can contain both disk pools and disk groups.

Disk Pool Restrictions

CAUTION: If you downgrade the RAID controller module firmware version of a storage array that is configured with a disk pool to a firmware version that does not support disk pools, the virtual disks are lost and the physical disks are treated as unaffiliated with a disk pool.

- All physical disk media types in a disk pool must be the same. Solid State Disks (SSDs) are not supported.
- You cannot change the segment size of the virtual disks in a disk pool.
- You cannot export a disk pool from a storage array or import the disk pool to a different storage array.

- You cannot change the RAID level of a disk pool. MD Storage Manager automatically configures disk pools as RAID level 6.
- All physical disk types in a disk pool must be the same.
- You can protect your disk pool with Self Encrypting Disk (SED), but the physical disk attributes must match. For example, SED-enabled physical disks cannot be mixed with SED-capable physical disks. You can mix SED-capable and non SED-capable physical disks, but the encryption abilities of the SED physical disks cannot be used.

Creating A Disk Pool Manually

You can use the unconfigured capacity in a storage array to create a disk pool.

NOTE: Ensure that you have created virtual disks before you create a disk pool.

To create a disk pool:

1. Select the **Storage & Copy Services** tab.
2. Select the unconfigured capacity node.
3. From the menu bar, select **Storage** → **Disk Pool** → **Create.** Alternatively, right-click unconfigured capacity in the object tree and select **Create Disk Pool**.

   The **Create Disk Pool** window is displayed.
4. Type a name for the disk pool in **Disk pool name**.
5. Select one of these options in **Physical disk security**:

   – **Only security-capable physical disks** — To create a secure disk pool from security capable physical disks.

   NOTE: The **Only security-capable physical disks** option is available only when a security key is set up for the storage array.

   – **Any available physical disks** — To create a disk pool comprised of physical disks that may or may not be security capable or are a mix of security levels.

   NOTE: You can mix Self Encrypting Disk (SED)-capable and non SED-capable physical disks. However, the encryption abilities of the SED-capable physical disks cannot be used, as the physical disk attributes do not match.

   Based on the physical disk type and physical disk security type that you have selected, the **Disk pool candidates** table shows one or more disk pool configurations.
6. Locate the **Secure Enable?** column in the **Disk pool candidates** table and select the disk pool that you want to secure.

   NOTE: You can click **View Physical Disks** to view the details of the physical disks that comprise the selected disk pool configuration.
7. To send alert notifications when the usable capacity of the disk pool is reaching a specified percentage, perform the following steps:
   a) Click **View notification settings**.
   b) Select the check box corresponding to a critical warning notification.

      You also can select the check box corresponding to an early warning notification. The early warning notification is available only after you select the critical warning notification.
   c) Select or type a value to specify a percentage of usable capacity.

      When the configured (allocated) capacity in the disk pool reaches the specified percentage, an alert notification in the form of emails and SNMP trap messages are sent to the destination addresses that are specified in the **Configure Alerts** dialog. For more information about how to specify the destination addresses, see Configuring Alert Notifications.
8. Click **Create**.

Automatically Managing The Unconfigured Capacity In Disk Pools

The MD Storage Manager can detect the unconfigured capacity in a storage array. When the unconfigured capacity is detected, the MD Storage Manager prompts you to create one or more disk pools, or add the unconfigured capacity to an existing disk pool, or both. By default, the **Automatic Configuration** dialog is displayed when one of these conditions are true:

- The AMW is opened to manage a storage array, disk pools do not exist in the storage array, and there are enough similar physical disks to create a new disk pool.
- New physical disks are added to a storage array that has at least one disk pool. If there are enough eligible physical disks available, you can create a disk pool of different physical disk types than the existing disk pool.

  NOTE: If you do not want the **Automatic Configuration** dialog to be displayed again when unconfigured capacity is detected, you can select **Do not display again**. If you later want this dialog to be displayed again when unconfigured capacity is detected, you can select **Storage Array → Preferences** in the AMW to reset your preferences. If you do not want to reset the preferences, but do want to invoke the **Automatic Configuration** dialog, select **Storage Array → Configuration → Disk Pools.**

Each physical disk in a disk pool must be of the same physical disk type and physical disk media type and have similar capacity. If there are a sufficient number of physical disks of those types, the MD Storage Manager prompts you to create a single disk pool. If the unconfigured capacity consists of different physical disk types, the MD Storage Manager prompts you to create multiple disk pools.

If a disk pool is already defined in the storage array, and you add new physical disks of the same physical disk type as the disk pool, the MD Storage Manager prompts you to add the physical disks to the existing disk pool. If the new physical disks are of different physical disk types, the MD Storage Manager prompts you to add the physical disks of the same physical disk type to the existing disk pool, and to use the other physical disk types to create different disk pools.

NOTE: If there are multiple disk pools of the same physical disk type, a message is displayed indicating that the MD Storage Manager cannot recommend the physical disks for a disk pool automatically. However, you can manually add the physical disks to an existing disk pool. You can click **No** to close the **Automatic Configuration** dialog and, from the AMW, select **Storage Array → Disk Pool → Add Physical disks (Capacity).**

If additional physical disks are added to the storage array when the **Automatic Configuration** dialog is open, you can click **Update** to detect the additional physical disks. As a best practice, add all of the physical disks to a storage array at the same time. This action enables the MD Storage Manager to recommend the best options for using the unconfigured capacity.

You can review the options, and click **Yes** in the **Automatic Configuration** dialog to create one or more disk pools, or to add the unconfigured capacity to an existing disk pool, or both. If you click **Yes**, you also can create multiple equal-capacity virtual disks after the disk pool is created.

If you choose not to create the recommended disk pools, or not to add the unconfigured capacity to a disk pool, click **No** to close the **Automatic Configuration** dialog. You can then manually configure the disk pools by selecting **Storage Array → Disk Pool → Create from the AMW.**

Locating Physical Disks In A Disk Pool

You can use the **Blink** option to physically locate and identify all of the physical disks that comprise a selected disk pool.

To locate a disk pool:

1. Select the **Storage & Copy Services** tab.
2. Select the disk pool in the Tree view or the Table view.
3. From the menu bar, select **Storage → Disk Pool → Blink.**

   The LEDs on each physical disk that make up the selected disk pool blink.

4. Locate the physical disks in the disk pool, and click **OK**.

   The LEDs stop blinking.

   NOTE: If the LEDs for the disk pool do not stop blinking, from the AMW, select **Hardware → Blink → Stop All Indications.**

5. Click **OK**.

Renaming A Disk Pool

Use the **Rename** option to change the name of a disk pool when the current name is no longer meaningful.

Keep these guidelines in mind when you rename a disk pool:

- A disk pool name can consist of letters, numbers, and the special characters underscore (_), hyphen (-), and pound (#). If you choose any other characters, an error message is displayed. You are prompted to choose another name.
- Limit the name to 30 characters.
- Use a unique, meaningful name that is easy to understand and remember.
- Do not use arbitrary names or names that may quickly lose their meaning in the future.
- If you choose a disk pool name that is already in use, an error message is displayed. You are prompted to choose another name.

To configure alert notifications for a disk pool:

1. In AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Rename.** Alternatively, right-click on the disk pool and select **Rename**.

   The **Rename Disk Pool** dialog is displayed.
4. Type a new name in **Disk pool name**.
5. Click **OK**.

Configuring Alert Notifications For A Disk Pool

You can configure the MD storage manager to send alert notifications when the unconfigured (free) capacity of a disk pool is reaching a specified percentage. You can modify the alert notification settings after creating a disk pool.

To configure alert notifications for a disk pool:

1. In AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Change → Settings.**

   The **Change Disk Pool Settings** dialog is displayed.
4. In the **Change Warning Thresholds** area, select the check box corresponding to a critical warning notification.

   You also can select the check box corresponding to an early warning notification.

   NOTE: The early warning notification is available only after you select the critical warning notification.
5. Select or type a value to specify a percentage of usable capacity.

   When the unconfigured (free) capacity in the disk pool reaches the specified percentage, an alert notification in the form of e-mail messages and SNMP trap messages are sent to the destination addresses that are specified in the **Configure Alerts** dialog. For more information about how to specify the destination addresses, see Configuring Alert Notifications.

6. Click **OK**.

Adding Unassigned Physical Disks To A Disk Pool

Use the **Add Physical Disks (Capacity)** option to increase the free capacity of an existing disk pool by adding unassigned physical disks. After you add unassigned physical disks to a disk pool, the data in each virtual disk of the disk pool is redistributed to include the additional physical disks.

NOTE: Keep these guidelines in mind when you add physical disks to a disk pool:

- The status of the disk pool must be Optimal before you can add unassigned physical disks.
- You can add a maximum of 12 physical disks to an existing disk pool. However, the disk pool cannot contain more physical disks than the maximum limit for a storage array.
- You can add only unassigned physical disks with an Optimal status to a disk pool.
- The data in the virtual disks remains accessible during this operation.

To add unassigned physical disks a disk pool:

1. In AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Add Physical Disks (Capacity).**
   The **Add Physical Disks** dialog is displayed. You can view information about:
   - The disk pool in the **Disk Pool Information** area.
   - The unassigned physical disks that can be added to the disk pool in the **Select physical disks for addition** area.

     NOTE: The RAID controller module firmware arranges the unassigned physical disk options with the best options listed at the top in the **Select physical disks for addition** area.
4. Select one or more physical disks in the **Select physical disks for addition** area.
   The total free capacity that will be added to the disk pool is displayed in the **Total usable capacity selected** field.
5. Click **Add**.

Configuring The Preservation Capacity Of A Disk Pool

The preservation capacity in a disk pool is reserved for data reconstruction operations in case of physical disk failures.

To configure the preservation capacity in a disk pool:

1. In AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Change → Settings.**
   The **Change Disk Pool Settings** dialog is displayed.
4. In the **Preservation Capacity** area of the **Physical disks dedicated to preservation capacity** box, type or select a number of physical disks.
   The preservation capacity of the disk pool is dependent on the number of physical disks in the disk pool.
5. Click **OK**.

Changing The Modification Priority Of A Disk Pool

Use the **Modification Priority** option to specify the priority levels for modification operations in a disk pool relative to the system performance.

NOTE: Selecting higher priority for modification operations in a disk pool can slow the system performance.

The following are the priority levels to modify a disk pool:

- **Degraded Reconstruction Priority** — The degraded reconstruction priority level determines the priority of the data reconstruction operation when a single physical disk fails in a disk pool.
- **Critical Reconstruction Priority** — The critical reconstruction priority level determines the priority of the data reconstruction operation when at least two physical disks fail in a disk pool.
- **Background Operation Priority** — The background operation priority level determines the priority of the disk pool background operations, such as Virtual Disk Expansion (VDE) and Instant Availability Format (IAF).

To configure alert notifications for a disk pool:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Change → Settings.**

   The **Change Disk Pool Settings** dialog is displayed.
4. In the **Modification Priorities** area, move the slider bars to select a priority level.

   You can choose a priority level for:

   - Degraded reconstruction
   - Critical reconstruction
   - Background operation

   You can select one of the following priority levels:

   - **lowest**
   - **low**
   - **medium**
   - **high**
   - **highest**

   The higher the priority level, the larger is the impact on host I/O and system performance.

Changing The RAID Controller Module Ownership Of A Disk Pool

You can change the RAID controller module ownership of a disk pool to specify which RAID controller module must own all of the virtual disks in the disk pool.

Changing the RAID controller module ownership at the disk pool level causes each virtual disk in that disk pool to transfer to the other RAID controller module and use a new I/O path. If you do not want to set each virtual disk to the new path, change the RAID controller module ownership at the virtual disk level instead of the disk pool level.

CAUTION: Possible loss of data access – If you change the RAID controller module ownership while an application is accessing the virtual disks in the disk pool, it may result in I/O errors. Make sure that the application is not accessing the virtual disks, and there is a multi-path driver installed on the hosts before you perform this procedure.

To RAID controller module ownership of a disk pool:

1. In AMW, select the **Storage & Copy Services** tab.
2. Select the disk pool.
3. From the menu bar, select **Storage → Disk Pool → Change → Ownership/Preferred Path.**
4. Select the RAID controller module.
5. Click **Yes**.

Checking Data Consistency

Use the **Check Consistency** option to check the consistency on a selected disk pool or disk group.

Use this option only when instructed by the Recovery Guru.

CAUTION: Use this option only under the guidance of your Technical Support representative.

Keep these important guidelines in mind before you check data consistency:

- Disk pools are configured only as RAID Level 6.
- You cannot use this option on RAID Level 0 disk groups that have no consistency.
- If you use this option on a RAID Level 1 disk group, the consistency check compares the data on the replicated physical disks.
- If you perform this operation on a RAID Level 5 or RAID Level 6 disk group, the check inspects the parity information that is striped across the physical disks. The information about RAID Level 6 applies also to disk pools.
- To successfully perform this operation, these conditions must be present:
  - The virtual disks in the disk pool or disk group must be in Optimal status.
  - The disk pool or disk group must have no virtual disk modification operations in progress.
  - You can perform this operation only on one disk pool or disk group at a time. However, you can perform a consistency check on selected virtual disks during a media scan operation. You can enable a media scan consistency check on one or more virtual disks in the storage array.

To check data consistency:

1. Select the **Storage & Copy Services** tab.
2. Select the disk pool or disk group that you want to check.
3. Select one of the following from the menu bar:
   - **Storage → Disk Group → Advanced → Check Redundancy**
   - **Storage → Disk Pool → Advanced → Check Redundancy**
4. Click **Yes**.
5. Click **Start**.

   The check consistency operation starts and the **Check Consistency** dialog is displayed. The virtual disks in the disk pool or disk group are sequentially scanned, starting from the top of the table in the virtual disk dialog. The following actions occur as each virtual disk is scanned:

   - The virtual disk is selected in the virtual disk table.

     CAUTION: Possible loss of data access – A parity error is potentially serious and could cause a permanent loss of data.
   - The status of the consistency check is shown in the **Associated Status** column.

Deleting A Disk Pool

Use the **Delete** option to delete a disk pool and all of the virtual disks in the disk pool. When a disk pool is deleted, the physical disks that were associated with the disk pool change to the **Unassigned** state. This process creates more unconfigured capacity in the storage array, which you can reconfigure to meet your storage needs.

CAUTION: Possible loss of data access – Deleting a disk pool causes loss of all data on the virtual disks in the disk pool. Before performing this operation, back up the data on all of the virtual disks in the disk pool, stop all input/output (I/O), and unmount any file systems on the virtual disk.

Keep these guidelines in mind before you delete a disk pool:

- If you delete a disk pool that contains a snapshot repository virtual disk, you must delete the base virtual disk before you delete the associated snapshot virtual disk.
- The capacity from the physical disks that were previously associated with the deleted disk pool is added to either of these nodes:
  - An existing Unconfigured Capacity node.
  - A new Unconfigured Capacity node if one did not exist previously.
- You cannot delete a disk pool that has any of these conditions:
  - The disk pool contains a repository virtual disk, such as a snapshot group repository virtual disk, a replication repository virtual disk, or a Consistency Group member repository virtual disk. You must delete the logical component that has the associated repository virtual disk in the disk pool before you can delete the disk pool.

    NOTE: You cannot delete a repository virtual disk if the base virtual disk is in a different disk pool and you have not requested to delete that disk pool at the same time.
  - The disk pool contains a base virtual disk or a target virtual disk participating in a virtual disk copy operation with the status of In Progress.

To delete a disk pool:

1. Select the **Storage & Copy Services** tab.
2. Select one or more disk pools.
3. From the menu bar, select **Storage → Disk Pool → Delete.**

   The **Confirm Delete Disk Pool** dialog is displayed.
4. Type `yes` to confirm, and click **OK**.

   The **Delete Disk Pool - Progress** dialog is displayed while all of the virtual disks in the disk pool are being deleted.

Viewing Storage Array Logical Components And Associated Physical Components

You can view the logical components (virtual disks , disk pools, and disk groups) in a storage array, and then view the physical components (RAID controller modules, RAID enclosures, physical disks, and expansion enclosures) that are associated with a specific logical component.

1. To view the components, select the **Storage & Copy Services** tab.

   The object tree is displayed on the left, and the Properties pane is displayed on the right. The object tree provides a view of the components in the storage array in a tree structure. The components shown include the disk pools, the disk groups, the virtual disks, the free capacity nodes, and any unconfigured capacity for the storage array. The Properties pane displays detailed information about the component that is selected in the object tree.
2. To view the physical components that are associated with a component, perform one of these actions:

– Right-click a component, and select **View Associated Physical Components**.
– Select a component, and click **View Associated Physical Components** in the Properties pane.
– Select a component, and from the menu bar, select **Storage → Disk Pool → View Associated Physical Components.**

The associated physical components are displayed with a green triangle on top of them.

Secure Disk Pools

You can create a secure disk pool from security capable physical disks. The physical disks in a secure disk pool become security enabled. Read access from and write access to the physical disks is only available through a RAID controller module that is configured with the correct security key.

CAUTION: Possible loss of data access – When a disk pool is secured, the only way to remove security is to delete disk pool. Deleting the disk pool deletes all of the data in the virtual disks that comprise the disk pool.

Whenever the power is turned off and turned on again, all of the security-enabled physical disks change to Security Locked status. In this status, the data is inaccessible until the correct security key is provided by a RAID controller module. You can view the Physical Disk Security status of any disk pool in the storage array from the **Disk Pool Properties** dialog. The following status information is reported:

- Security Capable
- Secure

The following table provides a description of the security properties status of a disk pool.

Table 6. Security Properties Status of a Disk Pool

	Security Capable – Yes	Security Capable – No
Secure – Yes	The disk pool is composed of all SED physical disks and is in Secure status.	Not applicable. Only SED physical disks can be in Secure status.
Secure – No	The disk pool is composed of all SED physical disks and is in Non-Secure status.	The disk pool is not entirely composed of SED physical disks.

The **Secure Physical Disks** option is displayed in the **Disk Pool** menu. The **Secure Physical Disks** option is active if these conditions are true:

- The selected storage array is not security enabled but is comprised entirely of security capable physical disks.
- The storage array contains no snapshot copy base virtual disks or snapshot repository virtual disks.
- The disk pool is in Optimal status.
- A security key is set up for the storage array.

The **Secure Physical Disks** option is inactive if the above conditions are not true. The **Secure Physical Disks** option is inactive with a check mark to the left if the disk pool is already security enabled.

The **Create a secure disk** pool option is displayed in the **Create Disk Pool - Disk Pool Name and Physical Disk Selection** dialog. The **Create a secure disk pool** option is active only when the following conditions are met:

- The Physical Disk Security feature is activated.
- A security key is installed in the storage array.
- At least one security capable physical disk is installed in the storage array.
- All of the physical disks that you selected on the **Hardware** tab are security capable physical disks.

Changing Capacity On Existing Thin Virtual Disks

If the amount of space used by the host for read/write operations (sometimes called consumed capacity) exceeds the amount of physical capacity allocated on a standard virtual disk, the storage array cannot accommodate additional write requests until the physical capacity is increased. However, on a thin virtual disk, MD Storage Manager can automatically expand physical capacity of a thin virtual disk. You can also do it manually using **Storage → Virtual Disk → Increase Repository Capacity.** If you select the automatic expansion option, you can also set a maximum expansion capacity. The maximum expansion capacity enables you to limit the automatic growth of a virtual disk to an amount less than the defined virtual capacity.

NOTE: Since less than full capacity is allocated when you create a thin virtual disk, insufficient free capacity may exist when certain operations are performed, such as snapshot images and snapshot virtual disks. If this occurs, an alert threshold warning is displayed.

Creating A Thin Virtual Disk From A Disk Pool

NOTE: You can create thin virtual disks only from disk pools, not from disk groups.

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select a **Free Capacity** node in a disk pool.
   The thin virtual disks are listed under the **Disk Pools** node.
3. Select **Storage → Virtual Disk → Create → Virtual Disk.**
   The **Create Virtual Disk** window is displayed.
4. Select **Create thin virtual disk**.
5. Use the **New virtual capacity** box to indicate the virtual capacity of the new virtual disk and **Units** to indicate the specific capacity units to use (MB, GB, or TB).
   The minimum virtual capacity is 32 MB.
6. In the **Virtual disk name** box, enter a name for the virtual disk.
7. To map hosts to virtual disks, select **Map later**.
   The virtual disk is not assigned a LUN and is not accessible by any hosts until you go to the **Host Mappings** tab and assign a specific host and LUN to this virtual disk.
8. To use flash SSD cache, select **Use flash SSD cache**.
   Flash SSD cache provides read-only caching of user selected virtual disks on Solid State Disks (SSDs) to further improve the read performance of those virtual disks beyond conventional hard drives. This process of copying data transparently off hard drives and on to high-performance SSDs improves application I/O performance and response times.
   The **Use flash SSD cache** check box is disabled if:
   - no SSD cache is available
   - the disk pool is comprised of only SSD physical disks
   - the disk pool has different data service attributes from the SSD cache
   - you selected **Map Later**

   NOTE: When you are creating a thin virtual disk, the **Enable dynamic cache read prefetch** option is not available.
9. Click **Next**.
10. Do one of the following:

- Select **Use recommended capacity settings** and click **Next**.
- Select **Choose your own settings** and then select **Customize capacity settings (advanced)**. Click **Next** and go to step 11.

11. Use the **Preferred capacity** box to indicate the initial physical capacity of the virtual disk and the **Units** list to indicate the specific capacity units to use (MB, GB, or TB).

    NOTE: The physical capacity is the amount of physical disk space that is currently reserved for write requests. The physical capacity must be at least 4 GB in size, and cannot be larger than 256 GB.

    Based on the value that you entered in the previous step, the **Disk pool physical capacity candidates** table is populated with matching repository virtual disks. New repository candidates returned either matches the capacity you specify, or be rounded up to the closest 4 GB increment to make sure all the repository capacity is usable.

12. Select a repository from the table.

    Existing repositories are placed at the top of the list.

    NOTE: The benefit of reusing an existing repository is that you can avoid the initialization process that occurs when you create a new one.

13. If you want to change the repository expansion policy or warning threshold, click **View advanced repository settings**.

    - **Repository expansion policy** – Select either **Automatic** or **Manual**. When the consumed capacity gets close to the physical capacity, you can expand the physical capacity. The MD storage management software can automatically expand the physical capacity or you can do it manually. If you select **Automatic**, you also can set a maximum expansion capacity. The maximum expansion capacity allows you to limit the virtual disk's automatic growth below the virtual capacity. The value for the maximum expansion capacity must be a multiple of 4 GB.
    - **Warning threshold** – In the **Send alert when repository capacity reaches** field, enter a percentage. The MD Storage Manager sends an alert notification when the physical capacity reaches the full percentage.

14. Click **Finish**.

    The **Virtual Disk Successfully Created** window is displayed.

15. Click **OK**.

    If you want to create another virtual disk, click **Yes** on the **Do you want to create another virtual disk?** window. Perform any operating system modifications necessary on the application host so that the applications can use the virtual disk. For more information, see the MD Storage Manager Software Installation Guide for your operating system.

9

Using SSD Cache

The SSD Cache feature utilizes solid-state disk (SSD) physical disks to improve read-only performance in your storage array. SSD physical disks are logically grouped together to provide secondary cache for use with the primary cache in the RAID controller module memory.

Using SSD Cache improves application throughput and response times and delivers sustained performance improvement across diverse workloads, especially high-IOP workloads.

How SSD Cache Works

Following a host read, data is stored in DRAM and is copied from user-specified base virtual disks and stored on two internal RAID virtual disks (one per RAID controller module). These virtual disks are automatically created when you initially set up an SSD cache. Neither virtual disks is accessible for read-write operations and cannot be displayed or managed in the MD Storage Manager interface.

Simple virtual disk I/O mechanisms are used to move data to and from the SSD cache.

Storing the data on the SSD cache eliminates the need for repeated access to the base virtual disk. However, both SSD cache virtual disks count against the number of virtual disks supported on the storage array.

Benefits Of SSD Cache

Benefits of using the SSD Cache feature varies depending on your system configuration and network environment. However, workloads that typically benefit the most from using high-performance SSD cache include:

- Workloads where performance is limited by physical disk input/output processes (IOPs).
- Applications that generate a much higher percentage of physical disk reads versus physical disk writes.
- Repeated reads to the same and/or adjacent areas of the physical disk.
- Overall data accessed by an application is routinely less than potential SSD cache capacity. To determine whether this is the case, reviewing the number of virtual disks and sizes that are cached often yields a reliable estimate. The more virtual disks that are cached, the more likely it is that your application accesses more data capacity than can be configured in SSD cache.

Choosing SSD Cache Parameters

When you create an SSD cache, you can choose which I/O type best matches your applications:

- file system
- database
- web server

You also have the following options:

- capacity of the SSD cache from a list of possible candidates consisting of different counts of SSD physical disks.
- whether you want to enable SSD cache on all eligible virtual disks currently mapped to hosts
- whether to use SSD cache on existing virtual disks or when creating new virtual disks

SSD Cache Restrictions

The following restrictions apply to using SSD Cache feature:

- SSD cache is not supported on Snapshots (Legacy) virtual disks or PiT-based Snapshot images.
- If you import or export base virtual disks that are SSD cache enabled or disabled, the cached data is not imported or exported.
- Maximum usable SSD cache capacity on a storage array is dependent on the RAID controller module's primary cache capacity.
- You cannot remove the last physical disk in an SSD cache without first deleting the SSD cache.
- Only one SSD cache is supported per storage array.

Creating An SSD Cache

1. In the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - In the tree view, right click on **SSD Cache** and select **Create**.
   - From the menu bar, select **Storage → SSD Cache → Create.**

   The **Create SSD Cache** window is displayed.
3. Type a name for **SSD Cache name**.
4. Select an **I/O characteristic type** from one of the following:
   - **File System**
   - **Database**
   - **Web Server**
5. Select an appropriate option for **Data Assurance (DA)**.
6. Select an appropriate capacity from **SSD cache candidates**.

   A maximum of 5,120 GB of SSD cache is available in the usable capacity.

   NOTE: To view the physical disks that comprise the usable capacity, select the appropriate row under **SSD cache candidates**, and click **View Physical Disks**.
7. To enable SSD cache on existing virtual disks, select **Enable SSD cache on existing compatible virtual disks that are mapped to hosts**.
8. Click **Create**.

Viewing Physical Components Associated With An SSD Cache

To view the physical components associated with an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache. and do one of the following:
   - From the menu bar, select **Storage → SSD Cache → View Associated Physical Components.**
   - Right click on the SSD cache and select **View Associated Physical Components**.
   - In the Table view for the SSD cache, click **View Associated Physical Components**.

   The **View Associated Physical Components** window is displayed.
3. To view a physical disk type, select a disk type from **Physical Disk Type** and click **Show**.

To hide the displayed components, click **Hide**.

4. To view the components installed in the associated enclosure, click **View Enclosure Components**.

Locating Physical Disks In An SSD Cache

To locate physical disks in an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache and do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Blink.**
   - Right click on the SSD cache and select **Blink**.

   The LEDs on the physical disks comprising the SSD cache blink.
3. After locating the physical disks, click **OK**.
   The LEDs stop blinking.
4. If the LEDs for the disk group do not stop blinking, from the toolbar in AMW, select **Hardware → Blink → Stop All Indications.**
   If the LEDs successfully stop blinking, a confirmation message is displayed.
5. Click **OK**.

Adding Physical Disks To An SSD Cache

To add physical disks to an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache and do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Add Physical Disks (Capacity).**
   - Right click on the SSD cache and select **Add Physical Disks (Capacity)**.

   The **Add Physical Disks (Capacity)** window is displayed.
3. Select the physical disk that you want to add and click **Add**.
   The following are not listed in the **Add Physical Disks (Capacity)** window:
   - Physical disk(s) in a non-optimal state.
   - Physical disks which are not SSD physical disks.
   - Physical disks not compatible with the physical disks currently in the SSD cache.

Removing Physical Disks From An SSD Cache

To remove physical disks from an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache from which you want to remove physical disk(s).
3. Do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Remove Physical Disks (Capacity).**
   - Right click on the SSD cache and select **Remove Physical Disks (Capacity)**.

   The **Remove Physical Disks (Capacity)** window is displayed.

4. Select the physical disk that you want to add and click **Remove**.

Suspending Or Resuming SSD Caching

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache and do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Suspend.**
   - Right click on the SSD cache and select **Suspend**.

   In the Table view of the SSD cache, the **Status** is displayed as **Suspended**.
3. To resume SSD caching, do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Resume.**
   - Right click on the SSD cache and select **Resume**.

   In the Table view of the SSD cache, the **Status** is displayed as **Optimal**.

Changing I/O Type In An SSD Cache

To change the I/O type in an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Change I/O Type** and select an appropriate I/O type.
   - Right click on the SSD cache and select **Change I/O Type** and select an appropriate I/O type.

   The newly selected I/O characteristic type is displayed in the Table view for the selected SSD cache.

Renaming An SSD Cache

To rename an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache which you want to rename.
3. Do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Rename.**
   - Right click on the SSD cache and select **Rename**.

   The **Rename SSD Cache** window is displayed.
4. Type a new name for the SSD cache and click **OK**.

Deleting An SSD Cache

To delete an SSD cache:

1. In the AMW, select the **Storage & Copy Services** tab.
2. In the tree view, select the SSD cache and do one of the following:
   - From the menu bar, select **Storage → SSD Cache → Delete.**

– Right click on the SSD cache and select **Delete**.

The **Confirm Delete SSD Cache** window is displayed.

3. Type `yes` to confirm and click **Delete**.

Using The Performance Modeling Tool

The SSD Cache Performance Modeling tool helps you determine the performance improvement for SSD cache capacity when you run the performance modeling tool with a workload that has the same characteristics as what you run in production. The tool provides an estimate of performance using the following metrics: cache hit percentage and average response time. This tool shows actual performance for the physical SSD cache that you created.

To run the performance modeling tool:

1. In the AMW, select the **Storage & Copy Services** tab.
2. To access the performance modeling tool, highlight the **SSD Cache** node in the logical tree view.
3. Do one of the following:

   – From the menu bar, select **Storage → SSD Cache → Run Performance Modeling.**
   – Right click on the SSD cache and select **Run Performance Modeling**.

   The **SSD Cache Performance Modeling** window is displayed.
4. Review the information in the **Modeling Information** area of the **SSD Cache Performance Modeling** window.
5. Select one of the following options from **View results** to choose the format you want to view the results:

   – Response Time
   – Cache Hit %
6. Click **Start** to run the performance modeling tool.

   NOTE: Depending on the cache capacity and workload, it may take about 10 to 20 hours to fully populate the cache. There is valid information even after a run of a few minutes, but it takes a number of hours to obtain the most accurate predictions.

   NOTE: While the performance modeling tool is running, a progress bar is displayed in the main area of the window. You can close or minimize the window and the performance modeling continues to run. You can even close the MD Storage Manager and the performance modeling session continues to run.
7. To save the results of a performance modeling session, click **Save As** and save the data to a **.csv** file.

10

Premium Feature—Snapshot Virtual Disk

The following types of virtual disk snapshot premium features are supported on the MD storage array:

- Snapshot Virtual Disks using multiple point-in-time (PiT) groups
- Snapshot Virtual Disks (Legacy) using a separate repository for each snapshot

NOTE: This section describes the Snapshot Virtual Disk using PiT groups. If you are using the Snapshot Virtual Disk (Legacy) premium feature, see Premium Feature—Snapshot Virtual Disks (Legacy).

A snapshot image is a logical image of the content of an associated base virtual disk created at a specific point-in-time, often known as a *restore point*. This type of image is not directly readable or writable to a host since the snapshot image is used to save data from the base virtual disk only. To allow the host to access a copy of the data in a snapshot image, you must create a snapshot virtual disk. This snapshot virtual disk contains its own repository, which is used to save subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image.

Snapshot Virtual Disk Vs. Snapshot Virtual Disk (Legacy)

Any snapshot virtual disk image is a logical point-in-time image of a virtual disk. However, the Snapshot Virtual Disk premium feature differs from the Snapshot Virtual Disk (Legacy) premium feature in the following ways:

- The Snapshot Virtual Disk (Legacy) premium feature uses one data repository for each snapshot. The non-legacy Snapshot Virtual Disk premium feature uses one data repository for all snapshot images associated with a base virtual disk. As a result, only one write operation per virtual disk is required (not multiple sequential write operations).
- Snapshot Virtual Disk adds the concept of a snapshot group. Because there is only one repository for multiple snapshot images, the repository is associated with the snapshot group instead of with the snapshot image.
- Unlike a Snapshot Virtual Disk (Legacy), a snapshot image created with the Snapshot Virtual Disk premium feature is not directly read/write accessible by a host since the snapshot image is used only to save changed data for a base virtual disk.
- Either Snapshot Virtual Disk or Snapshot Virtual Disk (Legacy) images can be created from a base virtual disk, but not both.
- A snapshot image created using Snapshot Virtual Disk is created almost instantaneously and initially uses no disk space.
- Snapshot Virtual Disk images are always created inside snapshot groups; each snapshot image is associated with only one snapshot group.

Snapshot Images And Groups

Snapshot images can be created manually or automatically by establishing a schedule that defines the date and time you want the snapshot image created. The following objects can be included in a snapshot image:

- Standard virtual disks
- Thin provisioned virtual disks
- Consistency groups

To create a snapshot image, you must first create a snapshot group and reserve snapshot repository space for the virtual disk. The repository space is based on a percentage of the current virtual disk reserve.

You can delete the oldest snapshot image in a snapshot group either manually or you can automate the process by enabling the Auto-Delete setting for the snapshot group. When a snapshot image is deleted, its definition is removed from the system, and the space occupied by the snapshot image in the repository is released and made available for reuse within the snapshot group.

Snapshot Virtual Disk Read/Write Properties

A snapshot virtual disk can be designated as either read-only or read-write.

The following are the differences between the two:

- Read-Only snapshot virtual disks provide the host read access to a copy of the data contained in the snapshot image. However, the host cannot modify the snapshot image. A Read-Only snapshot virtual disk does require an associated repository.
- Read-Write snapshot virtual disks require an associated repository to provide the host write access to a copy of the data contained in the snapshot image. A Read-Write snapshot virtual disk requires its own repository to save any subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image. The snapshot is allocated from the storage pool from which the original snapshot image is allocated. All I/O writes to the snapshot image are redirected to the snapshot virtual disk repository that was allocated for saving data modifications. The data of the original snapshot image remains unchanged. For more information, see Understanding Snapshot Repositories.

Snapshot Groups And Consistency Groups

The Snapshot Virtual Disk premium feature supports the following types of snapshot groups:

- Snapshot groups — A snapshot group is a collection of point-in-time images of a single associated base virtual disk.
- Consistency groups — A consistency group is a group of virtual disks that you can manage as a single entity. Operations performed on a consistency group are performed simultaneously on all virtual disks in the group.

Snapshot Groups

The purpose of a snapshot group is to create a sequence of snapshot images on a given base virtual disk without impacting performance. You can set up a schedule for a snapshot group to automatically create a snapshot image at a specific time in the future or on a regular basis.

When creating a snapshot group, the following rules apply:

- Snapshot groups can be created with or without snapshot images.
- Each snapshot image is can be a member of only one snapshot group.
- Standard virtual disks and thin virtual disks are the only types of virtual disks that can contain a snapshot group. Non-standard virtual disks, such as snapshot virtual disks, cannot be used for snapshot groups.
- The base virtual disk can reside on either a disk group or a disk pool.
- Snapshot virtual disks and snapshot groups cannot exist on the same base virtual disk.

A snapshot group uses a repository to save all data for the snapshot images contained in the group. A snapshot image operation uses less disk space than a full physical copy because the data stored in the repository is only the data that has changed since the latest snapshot image.

A snapshot group is created initially with one repository virtual disk. The repository initially contains a small amount of data, then increases over time with subsequent data updates. You can increase the size of the repository by increasing the capacity of the repository, or add virtual disks to the repository.

Snapshot Consistency Groups

To perform the same snapshot image operations on multiple virtual disks, you can create a consistency group containing the virtual disks. Any operation performed on the consistency group is performed simultaneously on all of the virtual disks in that group, which creates consistent copies of data between each virtual disk. Consistency groups are commonly used to create, schedule or rollback virtual disks.

Each virtual disk belonging to a consistency group is referred to as a member virtual disk. When you add a virtual disk to a consistency group, the system automatically creates a new snapshot group that corresponds to this member virtual disk. You can set up a schedule for a consistency group to automatically create a snapshot image of each member virtual disk in the group at a specific time in the future or on a regular basis.

A consistency group, pools multiple virtual disks together, enabling you to take a snapshot of all virtual disks at the same point in time. This creates a synchronized snapshot of all the virtual disks and is especially suitable for applications that span multiple virtual disks. For example, a database application containing log files on one virtual disk and the database on another.

For consistency groups, the following apples:

- Consistency groups can be created initially with or without member virtual disks.
- Snapshot images can be created for a consistency group to enable consistent snapshot images between all member virtual disks.
- Consistency groups can be rolled back.
- A virtual disk can belong to multiple consistency groups.
- Only standard virtual disks and thin virtual disks can be included in a consistency group.
- Snapshots created using the Snapshot Virtual Disk (Legacy) premium feature cannot be included in a consistency group.
- A base virtual disk can reside on either a disk group or disk pool.

Understanding Snapshot Repositories

Repositories are system-created virtual disks used to hold write data for a snapshots, snapshot groups and consistency groups. During creation of either group or write-enabled snapshot virtual disk, an associated repository is automatically created. By default, one individual repository virtual disk is created for each group or snapshot. You can create the overall repository automatically using the default settings, or you can manually create the repository by defining specific capacity settings.

A snapshot virtual disk allows the host access to a copy of the data contained in a snapshot image. A snapshot image is not directly read or write accessible to the host and is used only to save data captured from the base virtual disk.

Consistency Group Repositories

A consistency group is made up of simultaneous snapshots of multiple virtual disks. Each virtual disk that belongs to a consistency group is referred to as a member virtual disk. When you add a virtual disk to a consistency group, the system automatically creates a new snapshot group that corresponds to this member virtual disk. A consistency group repository must be created for each member virtual disk in a consistency group in order to save data for all snapshot images in the group.

A consistency group snapshot image comprises multiple snapshot virtual disks. Its purpose is to provide host access to a snapshot image that has been taken for each member virtual disk at the same moment in time. A consistency group snapshot image is not directly read or write accessible to hosts; it is used only to save the data captured from the base virtual disk. The consistency group snapshot virtual disk can be designated as either read-only or read-write. Read-write consistency group snapshot virtual disks require a repository for each member virtual disk in order to save any

subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image. Each member repository is created when the consistency group snapshot virtual disk is created.

Ranking Repository Candidates

If you choose to create a repository manually, you can filter repository candidates for each member virtual disk by selecting either a percentage of the base virtual disk capacity or by specifying a preferred capacity in the **Snapshot Group Settings** window. Based on your selection, best repository candidates are displayed. Repository candidates shown contain both new and existing repository virtual disks residing on a disk group or disk pool.

For more information, see the online help topics.

Using Snapshot Consistency Groups With Remote Replication

Although a virtual disk can belong to multiple consistency groups, you must create separate consistency groups for snapshot images and Remote Replication.

Adding a base virtual disk containing a consistency group to a Remote Replication (non-legacy, asynchronous), the repository automatically purges the oldest snapshot image and sets the auto-delete limit to the maximum allowable snapshot limit for a consistency group.

Additionally, all member virtual disks belonging to both a snapshot consistency group and a Remote Replication group must belong to the same Remote Replication group.

Creating Snapshot Images

A snapshot image is a logical point-in-time image of the content of an associated base virtual disk. With snapshot images, you can create multiple copies of production data on your storage array much more quickly than full copies. Snapshot images track source changes from the time each snapshot image is created. You can create snapshot images for the following storage objects:

- Standard virtual disks
- Thin virtual disks
- Consistency groups

Before creating a snapshot image, consider these guidelines:

- If you attempt to create a snapshot image on a snapshot group and that snapshot group has reached its maximum number of snapshot images, you can retry creating snapshot images after doing one of the following:
  - Enable automatic deletion of snapshot images in the **Advanced Options** section of the Create wizard.
  - Manually delete one or more snapshot images from the snapshot group.
- If you attempt to create a snapshot image and either of the following conditions below are present, the creation may remain in **Pending** state:
  - The base virtual disk that contains this snapshot image is a member of an Remote Replication group.
  - The base virtual disk is currently synchronizing. When synchronization is complete, the snapshot image creation will complete.
- You cannot create a snapshot image on a failed virtual disk or on a snapshot group designated as **Reserved**.

Creating A Snapshot Image

You can create a snapshot image by either selecting a base virtual disk or by selecting an existing snapshot group.

To create a snapshot image from an existing base virtual disk:

1. From the AMW, select the base virtual disk you are copying and select **Copy Services** → **Snapshot Image** → **Create.**

   The **Select or Create a Snapshot Group** window is displayed.

2. Do one of the following:

   - If snapshot groups already exist on the base virtual disk or if the base virtual disk already has the maximum number of snapshot groups, the **An Existing Snapshot Group** radio button is selected by default. Go to step 3.
   - If the base virtual disk does not contain any snapshot groups, the following message is displayed: `There are no existing snapshot groups on this base virtual disk. Use the option below to create a new snapshot group`. You must create a snapshot group on the base virtual disk before you can proceed. Go to step 4.

3. If you want to create a snapshot image on an existing snapshot group:

   a) Select a snapshot group from the existing snapshot group table.

   NOTE: Ensure that you select a snapshot group that has not reached its maximum limit of snapshot images.

   b) Click **Finish** to automatically complete the snapshot image creation process and then go to step 5.

4. If you want to create a new snapshot group for the snapshot image, you must select how you wish to create the snapshot group repository. Do one of the following:

   - Select **Automatic** and click **Finish** to create the snapshot group repository with the default capacity settings. This is the recommended option. Go to step 5.
   - Select **Manual** and click **Next** to define the properties for the snapshot group repository. Then click **Finish** to continue with the snapshot image creation process. Go to step 5.

     NOTE: Use this option if you want to specify all of the customizable settings for the snapshot group repository. The **Manual** method is considered advanced. It is recommended that you fully understand physical disk consistency and optimal physical disk configurations before proceeding with the **Manual** method.

     NOTE: Make sure you have either existing repositories, enough free capacity nodes, or available unconfigured capacity for the storage array on which you are creating the snapshot group repository, otherwise you cannot create the repository and an error message is displayed.

5. Click **Finish**.

   The system takes a copy of the associated base virtual disk. The snapshot image and its properties are displayed in the navigation tree for the associated base virtual disk.

Canceling A Pending Snapshot Image

Use the **Cancel Pending Snapshot Image** option to cancel a snapshot image that was put in a **Pending** state when you attempted to create the snapshot image for either a snapshot group or a consistency group.

The snapshot image is in a **Pending** state due to the following concurrent conditions:

- The base virtual disk for a snapshot group or one or more member virtual disks of a consistency group that contains this snapshot image is a member of an asynchronous remote replication group.
- The virtual disk or virtual disks are currently in a synchronizing operation.

The snapshot image creation operation completes as soon as the synchronization operation is complete. To cancel the pending snapshot image creation before the synchronization operation completes, do the following:

1. From the AMW, select either the snapshot group or consistency group that contains the pending snapshot image.
2. Do one of the following:
   - **Copy Services → Snapshot Group → Advanced → Cancel Pending Snapshot Image.**
   - **Copy Services → Consistency Group → Advanced → Cancel Pending Consistency Group Snapshot Image.**

Deleting A Snapshot Image

Use the **Delete Snapshot Image** option to delete the oldest snapshot image from a snapshot group or consistency group.

After a snapshot image is deleted from a snapshot group, the system performs the following actions:

- Deletes the snapshot image from the storage array.
- Releases the repository's reserve space for reuse within the snapshot group.
- Disables all the associated snapshot virtual disks that exist for the deleted snapshot image.

For a consistency group you can delete:

- A single snapshot image.
- Multiple snapshot images that have the same sequence number and creation timestamp.

When a snapshot image(s) is deleted from a consistency group, the system performs the following actions:

- Deletes the snapshot image from the storage array.
- Releases the repository's reserve space for reuse within the consistency group.
- Moves any member virtual disk, associated with the deleted snapshot image(s), to a **Stopped** state.
- Disables the member snapshot virtual disks associated with the deleted snapshot image(s).

To delete the snapshot image, do the following:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot image that you want to delete from the snapshot group or consistency group and then select one of the following menu paths to delete the snapshot image:
   - **Copy Services → Snapshot Image → Delete.**
   - **Copy Services → Consistency Group → Consistency Group Snapshot Image → Delete.**

   The **Confirm Delete** window is displayed.
3. Type `yes` in the text box and then click **Delete** to delete the snapshot image.

Scheduling Snapshot Images

MD Storage Manager allows you to schedule regular snapshot image creation to enable file recovery, and scheduled backups. You can create a schedule for an existing snapshot group or consistency group or when you initially create a snapshot group or consistency group.

- You can set up a schedule for a snapshot group to automatically create a snapshot image at a specific time in the future or on a regular basis.
- You can set up a schedule for a consistency group to automatically create a snapshot image of each member virtual disk in the group at a specific time in the future or on a regular basis.

You can create a schedule that runs daily or weekly in which you select specific days of the week (Sunday through Saturday). To make scheduling easier, you can import an existing schedule for a snapshot group or consistency group. In addition, you can temporarily suspend scheduled snapshot image creation by disabling the schedule. When a schedule is disabled, the scheduled snapshot image creations do not occur.

Creating A Snapshot Schedule

The MD Storage Manager allows you to schedule regular snapshot image creation to enable file recovery, and scheduled backups. You can create a schedule when you initially create a snapshot group or consistency group, or you can add one later to an existing snapshot group or consistency group. After you create a snapshot image schedule, you can modify these and other schedule settings.

The following guidelines apply:

- Using a schedule can result in a large number of snapshot images, so be sure that you have sufficient repository capacity.
- Each snapshot group or consistency group can have only one schedule.
- Scheduled snapshot image creations do not occur when the storage array is offline or turned off.
- If you delete a snapshot group or consistency group that has a schedule, the schedule is also deleted.

The snapshot image creation operation completes as soon as the synchronization operation is complete. To cancel the pending snapshot image creation before the synchronization operation completes, do the following:

1. From the AMW, select either the snapshot group or consistency group that contains the pending snapshot image.
2. Do one of the following:
   - **Copy Services → Snapshot Group → Advanced → Cancel Pending Snapshot Image.**
   - **Copy Services → Consistency Group → Advanced → Cancel Pending Consistency Group Snapshot Image.**

   The **Schedule Settings** window is displayed.
3. Do one of the following:
   - If you want to use an existing schedule, click **Import settings from existing schedule**. The **Import Schedule** is displayed. Select the schedule you want to import from the **Existing schedules** table and then click **Import**.
   - If you want to create a new schedule, define the settings appropriately for the schedule. For more information on the settings, see the online help topics.
4. If you are creating the schedule for a snapshot group, select **Create the First Snapshot Image Now** to create a copy of the associated base virtual disk at the same time the schedule is created.
5. If you are creating the schedule for a consistency group, click **Enable snapshot image scheduling** to enable the scheduled snapshot image creation for the group.
6. Click **Finish** to create the schedule for the snapshot group or consistency group you selected.

   The system performs the following:
   - Creates the schedule for the snapshot group or consistency group and updates the **Properties** pane for the snapshot group or consistency group.
   - If you had selected **Create the First Snapshot Image Now**, the system takes a copy of the associated base virtual disk. Each subsequent snapshot image capture depends on the schedule you created.

Editing A Snapshot Schedule

Use the **Edit Snapshot Image Schedule** option to modify the schedule settings defined for a snapshot group or consistency group. You can also use the **Edit Snapshot Image Schedule** option to temporarily suspend scheduled

snapshot image creation by disabling the schedule. When a schedule is disabled, the scheduled snapshot image creations do not occur.

To edit a schedule:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot group or consistency group for which you want to edit a schedule.
3. Do one of the following:
   - **Copy Services → Snapshot Group → Edit Snapshot Image Schedule.**
   - **Copy Services → Consistency Group Snapshot Image → Create/Edit Schedule.**

   The **Edit Snapshot Image Schedule** window is displayed.
4. Do one of the following:
   - If you want to disable the schedule, de-select **Enable Snapshot Image Scheduling**.
   - If you want to use a different existing schedule, click **Import settings from existing schedule**. The **Import Schedule Settings** dialog is displayed. Select the new schedule you want to import from the **Existing schedules** table and then click **Import**.
   - If you want to edit the schedule, modify the schedule settings. For more information on the schedule settings, see the online help.
5. Click **OK** to apply your changes to the schedule for the snapshot group or consistency group you selected.

Performing Snapshot Rollbacks

You can rollback snapshot operations by either:

- Creating a snapshot virtual disk of a snapshot image, which allows you to retrieve deleted files from that snapshot virtual disk (the base virtual disk remains undisturbed).
- Restoring a snapshot image to the base virtual disk, which allows you to roll back the base virtual disk to a previous point-in-time.

NOTE: The host will have immediate access to the new-rolled-back base virtual disk, but the existing base virtual disk will not allow the host read-write access after the rollback is initiated. You can create a snapshot of the base virtual disk just before you start the rollback to save the pre-rollback base virtual disk for recovery purposes.

Snapshot images are useful any time you want to roll back to a known good data set at a specific point in time. For example, before performing a risky operation on a virtual disk, you can create a snapshot image to enable "undo" capability for the entire virtual disk. You can start a rollback from the following types of snapshot images:

- Snapshot image of a base virtual disk, which allows you to roll back the base virtual disk associated with a snapshot group to a previous state.
- Consistency group snapshot image, which allows you to roll back all or select member virtual disks of the consistency group to a previous state.

Snapshot Rollback Limitations

- The rollback operation does not change the content of the snapshot images that are associated with the base virtual disk.
- You cannot perform the following actions when a rollback operation is in progress:
  - Delete the snapshot image that is being used for the rollback.
  - Create a new snapshot image for a base virtual disk that is participating in a rollback operation.
  - Change the associated snapshot group's Repository-Full Policy.

- You cannot start a rollback operation when any of these operations are in progress in the storage array:
  - Expanding the capacity of a disk group.
  - Virtual disk Expansion (VDE) to increase the capacity of a virtual disk.
  - Migrating a disk group to a different RAID level.
  - Changing the segment size of a virtual disk.
- You cannot start a rollback operation if the base virtual disk is participating in a virtual disk copy.
- You cannot start a rollback operation if the base virtual disk is a secondary virtual disk in a remote replication. However, if the base virtual disk is the primary virtual disk in a remote replication, you can start a rollback operation. Additionally, you cannot perform a role reversal in a remote replication if the primary virtual disk is participating in a rollback operation.
- A rollback operation fails if any of the used capacity in the associated snapshot repository virtual disk has unreadable sectors.

NOTE: You also can use the command line interface (CLI) to start a rollback operation from multiple snapshot images concurrently, cancel a rollback operation, resume a rollback operation, modify the priority of a rollback operation, and view the progress of a rollback operation.

Starting A Snapshot Rollback

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select a snapshot image of a base virtual disk and then select **Copy Services → Snapshot Image → Rollback → Start.**
   - Select a consistency group snapshot image and then select **Copy Services → Consistency Group Snapshot Image → Rollback → Start.**

   Depending on your selection, either the **Confirm Rollback Snapshot Image** or the **Confirm Rollback Snapshot Image** window is displayed.
3. If you are starting the rollback operation from a consistency group snapshot image, select the member virtual disks from the member virtual disks table that you want to rollback; otherwise skip to step 4.
4. In the **Rollback priority** area, use the slider bar to set a priority for the rollback operation.
   - There are five priority rates available: lowest, low, medium, high, and highest.
   - If the priority is set at the lowest rate, I/O activity is prioritized and the rollback operation takes longer time to complete.
   - If the priority is set at the highest priority rate, the rollback operation is prioritized, but I/O activity for the storage array may be affected.
5. To confirm and start the rollback operation, type `yes` in the text box, and click **Rollback**.

   You can view the progress of the rollback operation in the **Properties** pane when you select the base virtual disk or the consistency group member virtual disk in the **Logical** pane.

Resuming A Snapshot Image Rollback

Use the **Resume Rollback** option to resume a rollback operation that is in a **Paused** state. The rollback operation is paused if an error occurs during the rollback operation.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select a snapshot image of either a base virtual disk or of a consistency group's member virtual disk and then select **Copy Services → Snapshot Image → Rollback → Resume.**

   The **Confirm Resume Rollback** window is displayed.

3. Click **Resume**.

   The following may occur depending on the error condition:

   - If the resume rollback operation is successful — You can view the progress of the rollback operation in the Properties pane when you select the base virtual disk or the consistency group member virtual disk in the Logical pane.
   - If the resume rollback operation is not successful — The rollback operation is paused again. The base virtual disk or member virtual disk displays Needs Attention icons, and the controller logs the event to the Major Event Log (MEL). You can follow the Recovery Guru procedure to correct the problem or contact your Technical Support representative.

Canceling A Snapshot Image Rollback

Use the **Cancel Rollback** option to cancel a rollback operation after it has been started. You can cancel an active rollback that is in progress (actively copying data), a pending rollback (in a pending queue awaiting resources to start), or a rollback that has been paused due to an error. When you cancel a rollback operation as it is in progress, the base virtual disk reverts to an unusable state and appears as failed in the MD Storage Manager. Therefore, consider canceling a rollback operation only when recovery options exist for restoring the content of the base virtual disk.

After you cancel a rollback operation, you must take one of the following actions:

- Reinitialize the content of the base virtual disk.
- Perform a new rollback operation to restore the base virtual disk (using either the same snapshot image that was used in the **Cancel Rollback** operation or a different snapshot image to perform the new rollback operation).

NOTE: If the snapshot group on which the snapshot image resides has one or more snapshot images that are automatically purged, the snapshot image used for the rollback operation may not be available for future rollbacks.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select a snapshot image of either a base virtual disk or of a consistency group's member virtual disk and then select **Copy Services → Snapshot Image → Rollback → Advanced → Cancel.**

   The **Confirm Cancel Rollback** window is displayed.
3. Click **Resume**.
4. Click **Yes** to cancel the rollback operation.
5. Type `yes` in the text box, and click **OK**.

Viewing The Progress Of A Snapshot Rollback

You can view the progress of the rollback operation in the **Properties** pane of the AMW when you select the base virtual disk or consistency group member virtual disk in the **Logical** pane.

When a rollback operation is in progress, the following information is displayed:

- The **Operation in Progress** bar at the bottom of the **Properties** pane.
- The time remaining.

The Rollback operation is a long-running operation. The **Operations in Progress** window displays all of the long-running operations that are currently running on the storage array. From this window, you can view the progress of the rollback operation for a snapshot image and its associated base virtual disk or consistency group member virtual disk.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the storage array for which you want to display the operations in progress.

   The **Operations in Progress** window is displayed.

3. To view the progress for operations that affect a base virtual disk or a consistency group snapshot image, click the triangle next to a base virtual disk or a consistency group snapshot image to expand or collapse it.
4. To change the interval for refreshing the display, use the spinner box in the lower-right corner of the window, and click **Update**.
5. To refresh the display immediately, click **Refresh Now**.

Changing Snapshot Rollback Priority

You can set the priority for a rollback operation. Higher priority allocates more system resources for the rollback operation and might affect the overall system performance.

You can change the rollback priority at any of these times:

- Before the rollback begins
- While the rollback operation has a status of In Progress

There are five priority rates available: lowest, low, medium, high, and highest.

- If the priority is set at the lowest rate, I/O activity is prioritized and the rollback operation takes longer time to complete.
- If the priority is set at the highest priority rate, the rollback operation is prioritized, but I/O activity for the storage array may be affected.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select a snapshot image of either a base virtual disk or of a consistency group's member virtual disk and then select **Copy Services → Snapshot Image → Rollback → Change Priority.**
   - Select a consistency group of either a base virtual disk or of a consistency group's member virtual disk and then select **Copy Services → Consistency Group Snapshot Image → Rollback → Change Priority.**

   The **Change Rollback Priority** window is displayed.
3. In the rollback priority area, use the slider bar to set a priority for the rollback operation.

   If you are changing the priority for a consistency group snapshot image, the priority setting is applied to all member virtual disks in the selected consistency group.
4. Click **Change** to apply your changes to the rollback priority.

Creating A Snapshot Group

A snapshot group is a sequence of point-in-time images of a single associated base virtual disk. A snapshot group uses a repository to save data for all snapshot images contained in the group. The repository is created at the same time the snapshot group is created.

Keep these guidelines in mind when creating a snapshot group:

- When a base virtual disk that contains a snapshot group is added to an asynchronous remote replication group, the system automatically changes the repository full policy to automatically purge the oldest snapshot image and sets the auto-delete limit to the maximum allowable snapshot limit for a snapshot group.
- If the base virtual disk resides on a standard disk group, the repository members for any associated snapshot group, can reside on either a standard disk group or a disk pool. If a base virtual disk resides on a disk pool, all repository members for any associated snapshot group must reside on the same disk pool as the base virtual disk.
- You cannot create a snapshot group on a failed virtual disk.
- If you attempt to create a snapshot image, that snapshot image creation operation might remain in a Pending state because of the following conditions:

- The base virtual disk that contains this snapshot image is a member of an asynchronous remote replication group.
- The base virtual disk is currently in a synchronizing operation. The snapshot image creation completes as soon as the synchronization operation is complete.

1. From the AMW, select the base virtual disk whose data you want to copy.
2. Select a snapshot image of a base virtual disk and then select **Copy Services → Snapshot Group → Create.**

   The **Snapshot Group Settings** window is displayed.
3. In the **Snapshot group name** field, enter a unique name (30 character maximum) that best describes the virtual disk selected for this group. For example, AccountingData.

   By default, the snapshot group name is shown in the name text box as:

   ```
   [base-virtual disk-name] - SG + sequence-number
   ```

   In this example, SG (snapshot group) is the appended suffix and sequence-number is the chronological number of the snapshot group relative to the base virtual disk.

   For example, if you create the first snapshot group for a base virtual disk called "Accounting", the default name of the snapshot group is "Accounting_SG_01". The default name of the next snapshot group you create based on "Accounting" is "Accounting_SG_02". There is a 30-character limit. After you reach this limit, you can no longer type in the text box. If the base virtual disk is 30 characters, the default name for the group uses the base virtual disk name truncated enough to add the suffix "SG" and the sequence string.
4. Select **Create the first Snapshot Image Now** to take the first copy of the associated base virtual disk at the same time the snapshot group is created.
5. Do one of the following to select how you want to create the snapshot group repository:

   - Select **Automatic** and click **Finish** to create the snapshot group repository with the default capacity settings. This option is the recommended one.
   - Select **Manual** and click **Next** to define the properties for the snapshot group repository; then click Finish to continue with the snapshot group creation process.

   NOTE: Use this option if you want to specify all of the customizable settings for the snapshot group repository. The Manual method is considered advanced and only those who understand physical disk redundancy and optimal physical disk configurations should use this method. See Creating The Snapshot Group Repository (Manually) for instructions on how to set the repository parameters.
6. Click **Finish**.

   The system performs the following actions:

   - The snapshot group and its properties under the individual virtual disk node for the associated base virtual disk are displayed in the navigation tree.
   - If **Create the first Snapshot Image Now** was selected, the system takes a copy of the associated base virtual disk and the **Snapshot Image Successfully Created** window is displayed.

Creating A Consistency Group Repository (Manually)

During the creation of a consistency group, a consistency group repository is created to store the data for all the snapshot images contained in the group. A consistency group's repository is created initially with one individual repository virtual disk.Each virtual disk that belongs to a consistency group is referred to as a member virtual disk. When you add a virtual disk to a consistency group, the system automatically creates a new snapshot group that corresponds to this member virtual disk. A consistency group repository must be created for every member virtual disk in the consistency group to save the data for all the snapshot images contained in the group.

The Manual method is considered advanced and only those who understand physical disk consistency, provisioning, and optimal physical disk configurations should use this method.

Keep these guidelines in mind when you name a consistency group:

- There is a minimum required capacity for a consistency group repository (depending on your configuration).
- When you define the capacity requirements for a repository, keep in mind any future requirements that you might have for other virtual disks in this disk group or disk pool. Make sure that you have enough capacity to meet your data storage needs, but you do not over allocate because you can quickly use up all the storage in your storage array.
- The list of repository candidates can contain both new and existing repository virtual disks. Existing repository virtual disks are left on the storage array by default when you delete a consistency group. Existing repository virtual disks are placed at the top of the list. The benefit to reusing an existing repository virtual disk is that you can avoid the initialization process that occurs when you create a new one.

To create a consistency group repository:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select **Copy Services→ Consistency Group→ Create.**
   The **Consistency Group Settings** window is displayed.
3. Select **Manual** and click **Next** to customize the repository candidate settings for the consistency group.
   The **Consistency Group Repository Settings - Manual** window is displayed.
4. Select how you want to filter the repository candidates for each member virtual disk in the consistency group, based on either a percentage of the base virtual disk capacity or by preferred capacity.
   For more information on the filter options, see the online help topics.
   The best repository candidate for each member virtual disk based on the selections you made is displayed.
5. Select **Edit individual repository candidates** if you want to edit repository candidates for the member virtual disks.
6. Select the repository, from the **Repository candidates** table, that you want to use for each member virtual disk in the consistency group.

   NOTE: Select a repository candidate that is closest to the capacity you specified.

   - The **Repository candidates** table shows both new and existing repositories that are capable of being used for each member virtual disk in the consistency group based on the value you specified for percentage or the value you specified for preferred capacity.
   - By default, the system displays the repositories for each member virtual disk of the consistency group using a value of 20% of the member virtual disk's capacity. It filters out undersized repository candidates, and those with different Data Service (DS) attributes. If appropriate candidates are not returned using these settings, you can click **Run Auto-Choose** to provide automatic candidate recommendations.
   - The **Difference** column shows the mathematical difference between your selected capacity and the actual capacity of the repository candidate. If the repository candidate is new, the system uses the exact capacity size that you specified and displays zero (0) in the **Difference** column.
7. To edit an individual repository candidate:
   a) Select the candidate from the **Repository candidates** table and click **Edit** to modify the capacity settings for the repository.
   b) Click **OK**.
8. Select **View advanced options** and then accept or change the following default settings as appropriate.
   For more information on the advanced settings, see the online help topics.
9. Click **Finish**.

Changing Snapshot Group Settings

Use the **Snapshot Group Change Settings** option to modify the auto-delete settings and the snapshot group repository settings that were configured when you created the snapshot group.

- **Auto-Delete Settings** — You can configure each snapshot group to keep the total number of snapshot images in the group at or below a user-defined maximum. When this option is enabled, the system automatically deletes the oldest snapshot image in the group, any time a new snapshot is created, to comply with the maximum number of snapshot images allowed for the group.
- **Snapshot Group Repository Settings** — You can define a maximum percentage for the snapshot group repository that determines when a warning is triggered when the capacity of a snapshot group repository reaches the defined percentage. In addition, you can specify which policy to use when the capacity of the snapshot group repository reaches its maximum defined percentage:
  - **Automatically purge oldest snapshot image** — The system automatically purges the oldest snapshot image in the snapshot group, which releases the repository's reserve space for reuse within the snapshot group.
- **Reject writes to base virtual disk**: When the repository reaches its maximum defined percentage, the system rejects any I/O write request to the base virtual disk that triggered the repository access.

1. From the AMW, select the **Storage & Copy Services** tab.
2. From the snapshot groups category node, select the snapshot group that you want to change and then select **Copy Services → Snapshot Group → Change Settings.**

   The **Change Snapshot Group Settings** window is displayed.
3. Change the snapshot group settings as required.
4. Click **OK** to apply your changes to the snapshot group.

Renaming A Snapshot Group

Use the **Rename Snapshot Group** option to change the name of the snapshot group when the current name is no longer meaningful or applicable.

Keep these guidelines in mind when you name a snapshot group:

- A name can consist of letters, numbers, and the special characters underscore (_), hyphen (-), and pound (#). If you choose any other characters, an error message is displayed. You are prompted to choose another name.
- Limit the name to 30 characters. Any leading and trailing spaces in the name are deleted.
- Use a unique, meaningful name that is easy to understand and remember.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.
- If you try to rename a snapshot group with a name that is already in use by another snapshot group, an error message is displayed, and you are prompted to choose another name for the group.

To rename a snapshot group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot group that you want to rename and then select **Copy Services → Snapshot Group → Rename.**

   The **Rename Snapshot Group** window is displayed.
3. Type a new name for the snapshot group and then click **Rename**.

Deleting A Snapshot Group

Use the **Delete Snapshot Group** option to delete a snapshot group.

The system performs the following actions when a snapshot group is deleted:

- Deletes all existing snapshot images from the snapshot group.
- Deletes the associated repository that exists for the snapshot group (if selected).
- Disables all the associated snapshot virtual disks that exist for the deleted snapshot images.

To delete the snapshot group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot group that you want to delete and then select **Copy Services → Snapshot Group → Delete.**
   The **Confirm Delete** window is displayed.
3. Select **Delete all repositories associated with this object?** if you want to delete the associated repository that exists for the snapshot group.
4. Type `yes` in the text box and then click **Delete** to delete the snapshot group.

Converting Snapshots (Legacy) To A Snapshot Group

Use the **Convert Snapshots (Legacy) to Snapshot Group** option a to convert a snapshot (Legacy) virtual disk and its associated repository to a snapshot group. The system performs the following actions for each converted snapshot (Legacy) virtual disk:

- Deletes the snapshot (Legacy) virtual disk definition and creates a new snapshot group (the new snapshot group is created empty, with no snapshot images).
- Converts the associated snapshot (Legacy) repository virtual disk to a snapshot group repository.
- Retains the same schedule (if a schedule has been defined) for the new snapshot group.
- Creates a read-only snapshot virtual disk with a **Paused** status. The new snapshot virtual disk inherits the World-Wide Name (WWN) and host mappings as the converted snapshot (Legacy) virtual disk.

NOTE: If the number of snapshots (Legacy) that exist for a given base virtual disk exceeds the maximum number of allowed snapshot groups per base virtual disk, then you must delete any excess snapshots (Legacy) before performing the conversion process.

Conversion Limitations

- The conversion process is performed on a given base virtual disk, and applies to all snapshots of a given base virtual disk.
- If there is an online virtual disk copy related snapshot (legacy) on the base virtual disk, you must delete the online virtual disk copy job before initiating the conversion process.
- Snapshot (Legacy) virtual disks and snapshot groups cannot exist on the same base virtual disk. Therefore, any snapshot (Legacy) virtual disk that you do not select for conversion is deleted from the storage array.

Converting Snapshots (Legacy)

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select either a base virtual disk or a snapshot (legacy) virtual disk and then select one of the following menu paths (depending your selection):

- **Storage → Virtual Disk → Convert Snapshot (Legacy) to Snapshot Group.**
- **Copy Services → Snapshot (Legacy) → Convert to Snapshot Group.**

3. In the **Available snapshot (legacy) virtual disks** table, select the snapshot (legacy) virtual disks that you want to convert and then click **Add** to add them to the **Snapshot (Legacy) virtual disks to convert** table.
4. Click **Convert** to convert the selected snapshots (Legacy) to snapshot groups.

   Snapshot (Legacy) virtual disks are no longer displayed in the left pane of the Logical view. Snapshot groups are added to the right pane under the appropriate Category nodes.If there are unselected snapshots (Legacy) (to be deleted), then the **Confirm Delete** window is displayed.
5. Type `yes` in the text box and then click **OK** to confirm that you want to perform the conversion operation.

Converting A Snapshot Virtual Disk To Read-Write

Use the **Convert Snapshot Virtual Disk to Read-Write** option a to convert a read-only snapshot virtual disk to a read-write snapshot virtual disk.

You can use the **Convert Snapshot Virtual Disk to Read-Write** option for these storage objects:

- Snapshot virtual disk
- Consistency group member's snapshot virtual disk

The conversion operation requires that a repository be provisioned to support write operations on the snapshot virtual disk.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select either a snapshot virtual disk or a consistency group member's snapshot virtual disk and then select **Copy Services → Snapshot Virtual disk → Convert to Read-Write.**
3. Select how you wish to create the repository for the Read-Write snapshot virtual disk. Do one of the following:

   - Select **Automatic** to create the snapshot virtual disk repository with the default capacity settings. This is the recommended option.
   - Select **Manual** to define the properties for the snapshot virtual disk repository. Use this option if you want to specify all of the customizable settings for the snapshot virtual disk repository. The **Manual** method is considered advanced and only those who understand physical disk consistency and optimal physical disk configurations should use this method. See the online help topics for more information on how to set the repository parameters.
4. Click **Convert** to convert the read-only snapshot virtual disk to read-write.

   The snapshot virtual disk or consistency group member's snapshot virtual disk table as read- write is displayed under the **Mode** column, and the **Repository** columns are now populated.

Viewing Associated Physical Components Of An Individual Repository Virtual Disk

You can use the **View Associated Physical Components** option to view the physical components (RAID controller modules, RAID enclosures, physical disks, and expansion enclosures) that are associated with an individual repository virtual disk for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Consistency group member snapshot virtual disk
- Asynchronous remote replicated pair

1. Select the **Storage & Copy Services** tab.
2. Select the storage object for which you want to view the associated physical components and then select **Individual Repository Virtual Disk → View Associated Physical Components.**

Creating A Consistency Group

A consistency group is simultaneous snapshots of multiple virtual disks, thus ensuring consistent copies of a group of virtual disks. Each virtual disk that belongs to a consistency group is referred to as a member virtual disk. When you add a virtual disk to a consistency group, the system automatically creates a new snapshot group that corresponds to this member virtual disk.

The following guidelines apply:

- If the base virtual disk resides on a standard disk group, the repository members for any associated consistency group, can reside on either a standard disk group or a disk pool. If a base virtual disk resides on a disk pool, all repository members for any associated consistency group must reside on the same disk pool as the base virtual disk.
- You cannot create a consistency group on a failed virtual disk.
- A consistency group contains one snapshot group for each virtual disk that is a member of the consistency group. You cannot individually manage a snapshot group that is associated with a consistency group. Instead you must perform the manage operations (create snapshot image, delete snapshot image or snapshot group, and rollback snapshot image) at the consistency group level.
- If you attempt to create a consistency group snapshot image, the operation might remain in a Pending state because of the following conditions:
  - The base virtual disk that contains this consistency group snapshot image is a member of an asynchronous remote replication group.
  - The base virtual disk is currently in a synchronizing operation. The consistency group snapshot image creation completes as soon as the synchronization operation is complete.

To create a consistency group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select **Copy Services → Consistency Group → Create.**

   The **Consistency Group Settings** window is displayed.
3. In the **Consistency group name** field, enter a unique name (30-character maximum) that best describes the member virtual disks that you want to add for this group.

   By default, the consistency group name is shown in the name text box as:

   ```
   CG + sequence-number
   ```

   In this example, CG (Consistency Group) is the prefix and sequence-number is the chronological number of the consistency group, and is incremented based on how many consistency groups currently exist.
4. Select if you want to add the member virtual disks to the consistency group now or later:
   - Select Add members now and then from the eligible member virtual disks, select the virtual disks that you want to add as members to the consistency group. If you choose this method, you must create a repository for each member of the consistency group. Go to step 5. You can click the **Select all** check box to add all the virtual disks displayed in the **Eligible virtual disks** table to the consistency group.
   - Select **Add members later** and then click **Finish** to create the consistency group without member virtual disks. Go to step 6.

   The **Eligible virtual disks** table shows only those virtual disks that are capable of being used in the consistency group. To be eligible to be a member of a consistency group, a virtual disk cannot be in a Failed state and must contain less than the maximum allowable number of associated snapshot groups.

5. Select how you want to create the repositories for each member in the consistency group.
   - Select **Automatic** and click **Finish** to create the repositories with the default capacity settings. This option is the recommended one.
   - Select **Manual** and click **Next** to define the capacity settings for the repositories; and then click **Finish** to continue with the consistency group creation process. You can click **Edit individual repository candidates** to manually edit a repository candidate for each member virtual disk.

     NOTE: Use this option if you want to specify all of the customizable settings for the repositories. The Manual method is considered advanced and only those who understand physical disk consistency and optimal physical disk configurations should use this method. See the online help topics for more information on setting the repository parameters.
6. Click **Finish**.

   In the navigation tree, the consistency group and its properties are displayed under the **Consistency Groups** node.

Creating A Consistency Group Repository (Manually)

During the creation of a consistency group, a consistency group repository is created to store the data for all the snapshot images contained in the group. A consistency group's repository is created initially with one individual repository virtual disk.Each virtual disk that belongs to a consistency group is referred to as a member virtual disk. When you add a virtual disk to a consistency group, the system automatically creates a new snapshot group that corresponds to this member virtual disk. A consistency group repository must be created for every member virtual disk in the consistency group to save the data for all the snapshot images contained in the group.

The Manual method is considered advanced and only those who understand physical disk consistency, provisioning, and optimal physical disk configurations should use this method.

Keep these guidelines in mind when you name a consistency group:

- There is a minimum required capacity for a consistency group repository (depending on your configuration).
- When you define the capacity requirements for a repository, keep in mind any future requirements that you might have for other virtual disks in this disk group or disk pool. Make sure that you have enough capacity to meet your data storage needs, but you do not over allocate because you can quickly use up all the storage in your storage array.
- The list of repository candidates can contain both new and existing repository virtual disks. Existing repository virtual disks are left on the storage array by default when you delete a consistency group. Existing repository virtual disks are placed at the top of the list. The benefit to reusing an existing repository virtual disk is that you can avoid the initialization process that occurs when you create a new one.

To create a consistency group repository:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select **Copy Services→ Consistency Group→ Create.**

   The **Consistency Group Settings** window is displayed.
3. Select **Manual** and click **Next** to customize the repository candidate settings for the consistency group.

   The **Consistency Group Repository Settings - Manual** window is displayed.
4. Select how you want to filter the repository candidates for each member virtual disk in the consistency group, based on either a percentage of the base virtual disk capacity or by preferred capacity.

   For more information on the filter options, see the online help topics.

   The best repository candidate for each member virtual disk based on the selections you made is displayed.
5. Select **Edit individual repository candidates** if you want to edit repository candidates for the member virtual disks.
6. Select the repository, from the **Repository candidates** table, that you want to use for each member virtual disk in the consistency group.

   NOTE: Select a repository candidate that is closest to the capacity you specified.

- The **Repository candidates** table shows both new and existing repositories that are capable of being used for each member virtual disk in the consistency group based on the value you specified for percentage or the value you specified for preferred capacity.
- By default, the system displays the repositories for each member virtual disk of the consistency group using a value of 20% of the member virtual disk's capacity. It filters out undersized repository candidates, and those with different Data Service (DS) attributes. If appropriate candidates are not returned using these settings, you can click **Run Auto-Choose** to provide automatic candidate recommendations.
- The **Difference** column shows the mathematical difference between your selected capacity and the actual capacity of the repository candidate. If the repository candidate is new, the system uses the exact capacity size that you specified and displays zero (0) in the **Difference** column.

7. To edit an individual repository candidate:
   a) Select the candidate from the **Repository candidates** table and click **Edit** to modify the capacity settings for the repository.
   b) Click **OK**.
8. Select **View advanced options** and then accept or change the following default settings as appropriate.
   For more information on the advanced settings, see the online help topics.
9. Click **Finish**.

Renaming A Consistency Group

Use the **Rename Consistency Group** option to change the name of the consistency group when the current name is no longer meaningful or applicable.

Keep these guidelines in mind when you name a consistency group:

- A name can consist of letters, numbers, and the special characters underscore (_), hyphen (-), and pound (#). If you choose any other characters, an error message is displayed. You are prompted to choose another name.
- Limit the name to 30 characters. Any leading and trailing spaces in the name are deleted.
- Use a unique, meaningful name that is easy to understand and remember.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.
- If you try to rename a consistency group with a name that is already in use by another consistency group, an error message is displayed, and you are prompted to choose another name for the group.

To rename a consistency group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the consistency group that you want to rename and then select **Copy Services → Consistency Group → Rename.**
   The **Rename Consistency Group** window is displayed.
3. Type a new name for the consistency group and then click **Rename**.

Deleting A Consistency Group

Use the **Delete Consistency Group** option to delete a consistency group.

The system deletes the following when a consistency group is deleted:

- All existing snapshot images from the consistency group.
- All existing snapshot virtual disks from the consistency group.
- All the associated snapshot images that exist for each member virtual disk in the consistency group.
- All the associated snapshot virtual disks that exist for each member virtual disk in the consistency group.
- All associated repositories that exist for each member virtual disk in the consistency group (if selected).

To delete a consistency group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the consistency group that you want to delete and then select **Copy Services → Consistency Group → Delete.**

   The **Confirm Delete** window is displayed.
3. Select **Delete all repositories associated with this consistency group** if you want to delete the associated repository that exists for the consistency group.
4. Type `yes` in the text box and then click **Delete** to delete the consistency group.

Changing The Settings Of A Consistency Group

Use the **Change Consistency Group Settings** option to modify the auto-delete settings and the consistency group repository settings that were configured when you created the consistency group.

- **Auto-Delete Settings** — You can configure each consistency group to keep the total number of snapshot images in the group at or below a user-defined maximum. When this option is enabled, the system automatically deletes the oldest snapshot image in the group, any time a new snapshot is created, to comply with the maximum number of snapshot images allowed for the group.
- **Consistency Group Repository Settings** — You can define a maximum percentage for the consistency group member repository that determines when a warning is triggered when the capacity of a consistency group member repository reaches the defined percentage. In addition, you can specify which policy to use when the capacity of the consistency group repository reaches its maximum defined percentage:
  - **Automatically purge oldest consistency image**— The system automatically purges the oldest consistency image in the consistency group, which releases the repository's reserve space for reuse within the consistency group.
- **Reject writes to base virtual disk**— When the repository reaches its maximum defined percentage, the system rejects any I/O write request to the base virtual disk that triggered the repository access.

1. From the AMW, select the **Storage & Copy Services** tab.
2. From the consistency groups category node, select the consistency group that you want to change and then select **Copy Services → Consistency Group → Change Settings.**

   The **Change Consistency Group Settings** window is displayed.
3. Change the consistency group settings as required.
4. Click **OK** to apply your changes to the consistency group.

Adding A Member Virtual Disk To A Consistency Group

Use the **Add Member Virtual Disks** option to add a new member virtual disk to an existing consistency group. When a new member is added to a consistency group, you must also add a repository virtual disk.

Standard virtual disks and thin virtual disks are the only type of virtual disks that can be used for a consistency group. Non-standard virtual disks, such as snapshot (legacy) virtual disks, cannot be used for consistency groups. The base virtual disk can reside on either a disk group or a disk pool.

If you decide to re-create the snapshot virtual disk or consistency group snapshot virtual disk, you must choose a snapshot image from the same base virtual disk.

The following guidelines apply:

- The Snapshot premium feature must be enabled on the storage array.
- To add a new member virtual disk, the consistency group must have has less than the maximum number of allowable virtual disks (as defined by your configuration).

- If the base virtual disk resides on a standard disk group, the repository members for any associated consistency group can reside on either a standard disk group or a disk pool. If a base virtual disk resides on a disk pool, the repository members for any associated consistency group must reside on the same disk pool as the base virtual disk.
- You cannot add a member virtual disk that is in a failed state.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select the base virtual disk that you want to add to the consistency group and then select **Storage → Virtual disk → Copy Services → Add to Consistency Group.** The **Select Consistency Group and Repository** window is displayed.
   - Select the consistency group to which you want to add member virtual disks and then select **Copy Services → Consistency Group → Add Member Virtual Disks.** The **Select Virtual Disks and Repositories** window is displayed.
3. Depending on your selection in step 2, do one of the following:
   - In the **Select Consistency Group and Repository** window, select the consistency group from the **Consistency groups** table, to which you want add the base virtual disk.
   - In the **Select Virtual Disks and Repositories**, select the member virtual disks from the eligible virtual disks table, that you want to add to the consistency group. The eligible virtual disks table shows only those virtual disks that are capable of being used in the consistency group. You can click the **Select all** check box to add all the virtual disks displayed in the **Eligible virtual disks** table to the consistency group.
4. Select how you wish to create the repository for the member virtual disk(s) you are adding to the consistency group:
   - Select **Automatic** and click **Finish** to create the repository with the default capacity settings. This option is the recommended one.
   - Select **Manual** and click **Next** to define the capacity settings for the repository and then click **Finish**. See the online help topics for more information on how to set the repository parameters.

   Use the Manual option if you want to specify all of the customizable settings for the repository. The Manual method is considered advanced and only those who understand physical disk redundancy and optimal physical disk configurations should use this method.

   The new member virtual disk(s) for the consistency group are displayed in the **Member Virtual Disks** table.

Removing A Member Virtual Disk From A Consistency Group

Use the **Remove Member Virtual Disks** option to remove a member virtual disk from an existing consistency group. When you remove a member virtual disk from a consistency group, the system automatically deletes the snapshot group associated with that member virtual disk. In addition, you can choose whether you want to delete any repositories associated with the member virtual disk.

To remove a member virtual disk from a consistency group:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select the base virtual disk that you want to remove from the consistency group and then select **Storage → Virtual disk → Copy Services → Remove From Consistency Group.**
   - Select the consistency group to which you want to add member virtual disks and then select **Copy Services → Consistency Group → Remove Member Virtual Disks.**
3. If you selected a base virtual disk that is a member of multiple consistency groups or if you selected a consistency group from which you want to remove member virtual disk, do one of the following:

– Select one or more consistency groups, from the **Consistency groups** table, that you want to remove the base virtual disk from and then click **Remove**.

NOTE: You can click the **Select all** check box to remove the virtual disk from all the consistency groups displayed in the table.

– Select the member virtual disks, from the **Member virtual disks** table, that you want to remove from the consistency group and then click **Remove**.

NOTE: You can click the **Select all** check box to remove all the virtual disks displayed in the table.

4. Select the **Delete all repositories associated with this member virtual disk** if you want to delete all associated repositories that exist for the member virtual disk(s) in the consistency group.
5. Type `yes` in the text box and then click **Delete** to delete the member virtual disk(s) from the consistency group.

   The system removes the member virtual disks from the consistency group; they will not be deleted.

Creating A Snapshot Virtual Disk Of A Snapshot Image

You create a snapshot virtual disk to provide host access to a snapshot image within a snapshot group. A read-write snapshot virtual disk has its own repository that is used to save any subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image.

The snapshot virtual disk can be designated as either read-only or read-write:

- A read-only snapshot virtual disk provides a host application with READ access to a copy of the data contained in the snapshot image, but without the ability to modify the snapshot image. A read-only snapshot virtual disk does not have an associated repository.
- A read-write snapshot virtual disk requires an associated repository to provide the host application with WRITE access to a copy of the data contained in the snapshot image.

Snapshot Virtual Disk Limitations

- You cannot create a snapshot virtual disk of a Failed base virtual disk.
- Snapshot repositories are fully resizeable. If you have the storage capacity you can increase the size of the snapshot repository to avoid a repository full message. Conversely, if you find that the snapshot repository is larger than you need, you can reduce its size to free up space that is needed by other logical virtual disks.
- If you create a snapshot virtual disk for a snapshot image and that snapshot image creation operation remains in a Pending state it is due to the following conditions:
  - The base virtual disk that contains this snapshot image is a member of an asynchronous remote replication group.
  - The base virtual disk is currently in a synchronizing operation. The snapshot image creation will complete as soon as the synchronization operation is complete.

Creating A Snapshot Virtual Disk

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select a base virtual disk, and then select **Copy Services → Snapshot Virtual disk → Create.** The **Select Existing Snapshot Image or New Snapshot Image** window is displayed.
   - Select a base virtual disk, and then select **Copy Services → Snapshot Image → Create Snapshot Virtual Disk.** The **Snapshot Virtual Disk Settings** window is displayed. Go to step 4.
3. If you selected a base virtual disk in step 1, choose the snapshot image for which you want to create a snapshot virtual disk. Do one of the following:

- Select **An existing snapshot image** and then select a snapshot image from the snapshot image table and click **Next**.
- Select **A new snapshot image (on an existing snapshot group)** and then a snapshot group from the existing snapshot group table and then click **Next**.

The **Snapshot Virtual Disk Settings** window is displayed.

4. In the **Snapshot virtual disk name** field, enter a unique name (30 character maximum) that best describes the virtual disk selected for this snapshot image, for example, AccountingData.

   By default, the snapshot virtual disk name is shown in the name text box as follows:

   ```
   [base-virtual disk-name] - SV + sequence-number
   ```

   In this example, SV (snapshot virtual disk) is the appended suffix and sequence-number is the chronological number of the snapshot virtual disk relative to the base virtual disk.

   For example, if you create the first snapshot virtual disk for a base virtual disk called "Accounting", the default name of the snapshot virtual disk is "Accounting_SV_01". The default name of the next snapshot virtual disk you create based on "Accounting" is "Accounting_SV_02".

   There is a 30-character limit. After you reach this limit, you can no longer type in the text box. If the base virtual disk is 30 characters, the default name for the group uses the base virtual disk name truncated enough to add the suffix "SV" and the sequence string.

5. In the **Map to host** drop-down, specify how you want to map the host to the snapshot virtual disk.

   - **Map Now to Default Group** – The virtual disk is automatically assigned a logical unit number (LUN) and is accessible by any hosts that are connected to the storage array.
   - **Map Later** – The virtual disk is not assigned a LUN and is not accessible by any hosts until you go to the **Host Mappings** tab and assign a specific host and LUN to this virtual disk.
   - **Select a specific host** – You can select a specific host or host group from the list. This option is available only if Storage Partitioning is enabled.

   NOTE: Make sure there are enough free LUNs on the host or host group that you selected to map to a snapshot virtual disk.

6. Select how to grant host access to the snapshot virtual disk. Do one of the following:

   - Select **Read Write** and go to step 7.
   - Select **Read Only** and click **Finish** to create the snapshot virtual disk. Go to step 8.

     NOTE: Repositories are not required for **Read Only** snapshot virtual disks.

   Keep these guidelines in mind when you grant host access to a snapshot virtual disk:

   - Each host has its own logical unit number (LUN) address space and allows the same LUN to be used by different host groups or hosts to access snapshot virtual disks in a storage array.
   - You can define one mapping for each snapshot virtual disk in the storage array.
   - Mappings are shared between controllers in the storage array.
   - The same LUN cannot be used twice by a host group or a host to access a snapshot virtual disk. You must use a unique LUN.
   - An access virtual disk mapping is not required for out-of-band storage arrays.

7. Choose how you want to create the repository for the Read-Write snapshot virtual disk. Do one of the following:

   - Select **Automatic** and click **Finish** to create the snapshot virtual disk repository with the default capacity settings. This option is the recommended one.
   - Select **Manual** and click **Next** to define the properties for the snapshot virtual disk repository. Then click **Finish** to continue with the snapshot virtual disk creation procedure.

Use this option if you want to specify all of the customizable settings for the snapshot virtual disk repository. The Manual method is considered advanced and only those who understand physical disk redundancy and optimal physical disk configurations should use this method.

8. Click **Finish**.

   The snapshot virtual disk and its properties under the individual virtual disk node for the associated base virtual disk is displayed in the navigation tree. The snapshot virtual disk is added as a new virtual disk that contains the snapshot image information, which is the data of the virtual disk at the particular time of snapshot image creation.

Creating A Snapshot Virtual Disk Repository

When you create a snapshot virtual disk that is designated as read-write, a snapshot virtual disk repository is created to provide the host application with write access to a copy of the data contained in the snapshot image. You can create the repository automatically using the default settings or you can manually create the repository by defining the capacity settings for the repository.

The following guidelines apply:

- There is a minimum required capacity for a snapshot group repository which depends on your configuration.
- When you define the capacity requirements for a repository, keep in mind any future requirements that you may have for other virtual disks in this disk group or disk pool. Make sure that you have enough capacity to meet your data storage needs without allocating too much capacity that takes up the storage in your system.
- The list of repository candidates can contain both new and existing repository virtual disks. Existing repository virtual disks are placed at the top of the list. The benefit of reusing an existing repository virtual disk is that you can avoid the initialization process that occurs when you create a new one.

To create a snapshot virtual disk repository:

1. From the **Snapshot Virtual Disk Settings** window, select **Manual** and click **Next** to define the properties for the snapshot virtual disk repository.

   The **Snapshot Virtual disk Repository Settings - Manual** window is displayed.

2. Choose how you want filter the repository candidates in the **Repository candidates** table, based on either a percentage of the base virtual disk capacity or by preferred capacity.

   For more information, see the online help topics.

   The repository candidates that you selected are displayed.

3. Select the repository, from the **Repository candidates** table, that you want to use for the snapshot virtual disk and select a repository candidate that is closest to the capacity you specified.

   – The **Repository candidates** table shows both new and existing repositories that are capable of being used for the snapshot virtual disk based on the value you specified for percentage or the value you specified for preferred capacity.

   – The **Difference** column shows the mathematical difference between your selected capacity and the actual capacity of the repository candidate. If the repository candidate is new, then the system uses the exact capacity size that you specified and displays zero (0) in the **Difference** column.

4. In the **% Full** box, define the value that determines when a warning is triggered when the capacity of a snapshot virtual disk repository reaches the defined percentage.
5. Click **Finish**.

Changing The Settings Of A Snapshot Virtual Disk

Use the **Change Snapshot Virtual Disk Settings** option to modify the repository settings that were configured when you created the snapshot virtual disk. You can modify the maximum percentage for the snapshot virtual disk repository to set a warning when the capacity of a snapshot virtual disk repository reaches the defined percentage.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select a base virtual disk, and then select **Copy Services → Snapshot Virtual disk → Change Settings.**
   The **Change Snapshot Virtual Disk Settings** window is displayed.
3. Modify the repository full settings as required.
4. Click **OK** to apply the changes.

Disabling A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Disable** option when you want to invalidate a snapshot copy or a consistency group snapshot virtual disk. If the snapshot virtual disk or consistency group snapshot virtual disk is designated as read-write, this option also allows you stop any further write activity to its associated snapshot repository virtual disk.

Use the **Disable** option if one of these conditions applies:

- You are finished with the snapshot virtual disk or consistency group snapshot virtual disk for the time being.
- You intend to re-create the snapshot virtual disk or consistency group snapshot virtual disk (that is designated as read-write) at a later time and want to retain the associated snapshot repository virtual disk so that it does not need to be created again.
- You want to maximize the storage array performance by stopping write activity to the snapshot repository virtual disk.

If you decide to re-create the snapshot virtual disk or consistency group snapshot virtual disk, you must choose a snapshot image from the same base virtual disk.

If you disable the snapshot virtual disk or consistency group snapshot virtual disk, the system performs the following actions:

- Retains the World-Wide Name (WWN) for the snapshot virtual disk or consistency group snapshot virtual disk.
- Retains the snapshot virtual disk or consistency group snapshot virtual disk's association with the same base virtual disk.
- Retains the snapshot virtual disk or consistency group snapshot virtual disk's associated repository (if the virtual disk is designated as read-write).
- Retains any host mapping and access (any read-write requests will fail).
- Removes the snapshot virtual disk or consistency group snapshot virtual disk's association with the current snapshot image.
- For a consistency group snapshot virtual disk, disables each member's snapshot virtual disk.

NOTE: If you are finished with the snapshot virtual disk or consistency group snapshot virtual disk and do not intend to re-create it at a later time, you must delete the virtual disk, instead of disabling it.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:
   - **Copy Services → Snapshot Virtual disk → Disable.** The **Confirm Disable Snapshot Virtual Disk** window is displayed.

– **Copy Services → Consistency Group Snapshot Virtual Disk → Disable.** The **Confirm Disable Consistency Group Snapshot Virtual Disk** window is displayed.

3. Type `yes` in the text box and then click **Disable** to disable the snapshot virtual disk.

   The snapshot virtual disk or consistency group snapshot virtual disk is displayed in the Logical pane with the Disabled Snapshot status icon. If you disabled a read-write snapshot virtual disk or consistency group snapshot virtual disk, its associated snapshot repository virtual disk does not change status. The write activity to the snapshot repository virtual disk stops until the snapshot virtual disk or consistency group snapshot virtual disk is re-created.

Re-creating A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Re-Create** option when you want to re-create a snapshot virtual disk or consistency group snapshot virtual disk that you previously disabled. Re-creating a snapshot virtual disk or consistency group snapshot virtual disk takes less time than creating a new one.

If you have a snapshot virtual disk or consistency group snapshot virtual disk that you no longer need, you can reuse it (and any associated snapshot repository virtual disk), instead of deleting it, to create a different snapshot virtual disk or consistency group snapshot virtual disk of the same base virtual disk. You can re-associate the snapshot virtual disk or consistency group snapshot virtual disk with the same snapshot image or a different snapshot image as long as the snapshot image is in the same base virtual disk.

NOTE: If the snapshot virtual disk or consistency group snapshot virtual disk is part of an online copy relationship, you cannot perform the **Re-create** option on the virtual disk.

Keep these important guidelines in mind when you re-create a snapshot virtual disk or consistency group virtual disk:

- The snapshot virtual disk or consistency group snapshot virtual disk must be in either an **Optimal** status or **Disabled** status.
- For consistency group snapshot virtual disk, all member snapshot virtual disks must be in a Disabled state before you can re-create the consistency group snapshot virtual disk.
- You cannot re-create an individual member snapshot virtual disk, you can re-create only the overall consistency group snapshot virtual disk.
- All write data on any associated snapshot repository virtual disk is deleted. Snapshot virtual disk or consistency group snapshot virtual disk parameters remain the same as the previously disabled virtual disk parameters. The original names for the snapshot virtual disk or consistency group snapshot virtual disk are retained. You can change these names after the re-create option completes.

To re-create a snapshot virtual disk or consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:

   – **Copy Services → Snapshot Virtual disk → Re-create.** The **Confirm Re-Create Snapshot Virtual Disk** window is displayed.

   – **Copy Services → Consistency Group Snapshot Virtual Disk → Re-create.** The **Confirm Re-Create Consistency Group Snapshot Virtual Disk** window is displayed.

3. Select whether to re-create the snapshot virtual disk or consistency group snapshot virtual disk using an existing snapshot image, or a new snapshot image and then click **Re-create**.

   The status of the snapshot virtual disk or consistency group snapshot virtual disk is changed from **Disabled** to **Optimal**.

Renaming A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Rename Snapshot Virtual Disk** option to change the name of a snapshot virtual disk or consistency group snapshot virtual disk when the current name is no longer meaningful or applicable.

Keep these guidelines in mind when you name a consistency group:

- Limit the name to 30 characters. Any leading and trailing spaces in the name are deleted.
- Use a unique, meaningful name that is easy to understand and remember.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.

NOTE: If you try to rename a snapshot virtual disk or consistency group snapshot virtual disk with a name that is already in use by another virtual disk, an error message is displayed, and you are prompted to choose another name.

To rename a snapshot virtual disk or consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:
   - **Copy Services → Snapshot Virtual disk → Rename.** The **Rename Snapshot Virtual Disk** window is displayed.
   - **Copy Services → Consistency Group Snapshot Virtual Disk → Rename.** The **Rename Consistency Group** window is displayed.
3. Type a new name for the snapshot virtual disk or consistency group snapshot virtual disk and then click **Rename**.

Deleting A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Delete Snapshot Virtual Disk** option to delete a snapshot virtual disk or consistency group snapshot virtual disk that is no longer needed for backup or software application testing purposes. You also can specify whether you want to delete the snapshot repository virtual disk associated with a read-write snapshot virtual disk or a read-write consistency group snapshot virtual disk or retain the snapshot repository virtual disk as an unmapped virtual disk.

When a snapshot virtual disk or consistency group snapshot virtual disk is deleted, the system performs the following actions:

- Deletes all member snapshot virtual disks (for a consistency group snapshot virtual disk).
- Removes all associated host mappings.

NOTE: Deleting a base virtual disk automatically deletes any associated snapshot virtual disk or consistency group snapshot virtual disk. You cannot delete a snapshot virtual disk that is in a virtual disk copy with a status of **In Progress**.

To rename a snapshot virtual disk or consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:
   - **Copy Services → Snapshot Virtual disk → Delete.** The **Confirm Delete Snapshot Virtual Disk** window is displayed.
   - **Copy Services → Consistency Group Snapshot Virtual Disk → Delete.** The **Confirm Delete Consistency Group Snapshot Virtual Disk** window is displayed.

3. If the snapshot virtual disk or the consistency group snapshot virtual disk is read-write, select the option to delete the associated repository.
4. Type `yes` in the text box and then click **Delete** to delete the snapshot virtual disk or consistency group snapshot virtual disk.

Creating A Consistency Group Snapshot Virtual Disk

A consistency group snapshot virtual disk comprises multiple snapshot virtual disks to provide host access to a snapshot image that has been taken for each selected member virtual disk at the same moment in time. The consistency group snapshot virtual disk can be designated as either read-only or read-write. Read-write consistency group snapshot virtual disks require a repository for each member virtual disk that you select in the wizard to save any subsequent modifications made by the host application to the base virtual disk without affecting the referenced snapshot image. Each member repository is created at the same time the consistency group snapshot virtual disk is created.

The following guidelines apply:

- The Snapshot premium feature must be enabled on the storage array.
- The consistency group must contain at least one member virtual disk before you can create a consistency group snapshot virtual disk.
- There is a maximum allowable limit to the number of snapshot images for a consistency group (depending on your configuration).
- You cannot create a snapshot virtual disk of a failed virtual disk.
- Snapshot virtual disk repositories are fully resizeable. If you have the storage capacity you can increase the size of the snapshot repository to avoid a repository full message. Conversely, if you find that the snapshot virtual disk repository is larger than you need, you can reduce its size to free up space that is needed by other logical virtual disks.

NOTE: If you attempt to create a snapshot virtual disk for a snapshot image and that snapshot image is in a pending snapshot image creation operation, it is due to the following conditions:

- The base virtual disk that contains this snapshot image is a member of an asynchronous remote replication group
- The base virtual disk is currently in a synchronizing operation. The snapshot image is created as soon as the synchronization operation is completed.

To create a consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Do one of the following:
   - Select a consistency group, and then select **Copy Services → Consistency Group → Create Consistency Group Snapshot Virtual Disk.** The **Select Existing Snapshot Image or New Snapshot Image** window is displayed. Go to step 3.
   - Select a consistency group snapshot image from the **Consistency Group Snapshot Images** table, and then select **Copy Services → Consistency Group Snapshot Image → Create Consistency Group Snapshot Virtual Disk.** The **Consistency Group Snapshot Virtual Disk Settings** window is displayed. Go to step 4.
3. If you selected a consistency group in step 2, select the consistency group snapshot image for which you want to create a snapshot virtual disk. Do one of the following:
   - Select **An existing snapshot image** and then select a snapshot image from the consistency group snapshot images table and click **Next**.
   - Select **A new snapshot image** and then a snapshot group from the existing snapshot group table and then click **Next**.

   The **Consistency Group Snapshot Virtual Disk Settings** window is displayed.

4. In the **Consistency group snapshot virtual disk name** field, enter a unique name (30 character maximum) that best describes the consistency group selected for this snapshot image. For example, AccountingData.

   By default, the consistency group snapshot virtual disk name is shown in the name text box as:

   ```
   [consistency-group-name] - SV + sequence-number
   ```

   where SV (snapshot virtual disk) is the appended suffix and sequence-number is the chronological number of the snapshot virtual disk relative to the consistency group.

   For example, if you create the first snapshot virtual disk for a consistency group called "Accounting", then the default name of the snapshot virtual disk is "Accounting_SV_01". The default name of the next snapshot virtual disk you create based on "Accounting" is "Accounting_SV_02".

   There is a 30-character limit. After you reach this limit, you can no longer type in the text box. If the consistency group name is 30 characters, then the default name for the group uses the base virtual disk name truncated enough to add the suffix "SV" and the sequence string.

5. In the **Map to host** drop-down, specify how you want to map the host for each snapshot virtual disk created for a selected member virtual disk.

   This map attribute is applied to every member virtual disk you select in the consistency group. For more information on the map attributes, see the online help topics.

   The following guidelines apply:

   - Each host has its own logical unit number (LUN) address space and will let the same LUN be used by different host groups or hosts to access snapshot virtual disks in a storage array.
   - You can define one mapping for each snapshot virtual disk in the storage array.
   - Mappings are shared between RAID controller modules in the storage array.
   - The same LUN cannot be used twice by a host group or a host to access a snapshot virtual disk. You must use a unique LUN.
   - An access virtual disk mapping is not required for out-of-band storage arrays.

6. Select how to grant host access to each selected member virtual disk's snapshot virtual disk. Do one of the following:

   - Select **Read/Write** to provide the host application with WRITE access to a copy of the data contained in the snapshot image. A Read-Write snapshot virtual disk requires an associated repository.
   - Select **Read Only** to provide a host application with READ access to a copy of the data contained in the snapshot image, but without the ability to modify the snapshot image. A Read-Only snapshot virtual disk does not have an associated repository.

7. Select each member virtual disk in the consistency group for which you want to create a snapshot virtual disk.

   You can click **Select all** to create a snapshot virtual disk for each member virtual disk displayed in the select members table.

8. If you selected **Read-Only host access** in step 6, you can skip this step and go to step 9.

   NOTE: Repositories are not required for Read-Only snapshot virtual disks.

9. Select how you want to create the snapshot virtual disk repositories for each member in the consistency group. Do one of the following:

   - Select **Automatic** and click **Finish** to create each snapshot virtual disk repository with the default capacity settings. This option is the recommended one.
   - Select **Manual** and click **Next** to define the properties for each snapshot virtual disk repository; then click **Finish** to continue with the snapshot virtual disk creation process. You can click **Edit individual repository candidates** to manually edit a repository candidate for each member virtual disk.

   Use this option if you want to specify all of the customizable settings for the snapshot virtual disk repository. The Manual method is considered advanced and only those who understand physical disk consistency and optimal physical disk configurations should use this method.

See the online help topics for more information on how to set the repository parameters.

The snapshot virtual disk and its properties for the associated consistency group are displayed in the navigation tree.

Creating A Consistency Group Snapshot Virtual Disk Repository (Manually)

During the creation of a consistency group snapshot virtual disk that is designated as read-write, the system requires a snapshot virtual disk repository for each member of the consistency group to provide the host application with WRITE access to a copy of the data contained in the snapshot image. You can create the repository automatically using the default settings or you can manually create the repository by defining the capacity settings for the repository.

You are initially creating an overall repository with one individual repository virtual disk. However, the overall repository can contain multiple repository virtual disks in the future for expansion purposes.

Use the **Consistency Group Snapshot Virtual Disk Repository Settings - Manual** option to manually define the capacity requirements for a consistency group snapshot virtual disk repository. The Manual method is considered advanced and only those who understand physical disk redundancy, provisioning, and optimal physical disk configurations should use this method.

The following guidelines apply:

- There is a minimum required capacity for a snapshot virtual disk repository (depending on your configuration).
- When you define the capacity requirements for a repository, keep in mind any future requirements that you may have for other virtual disks in this disk group or disk pool. Make sure that you have enough capacity to meet your data storage needs, but you do not over allocate because you can quickly use up all the storage in your storage array.
- The list of repository candidates can contain both new and existing repository virtual disks. Existing repository virtual disks are left on the storage array by default when you delete a consistency group snapshot virtual disk. Existing repository virtual disks are placed at the top of the list. The benefit to reusing an existing repository virtual disk is that you can avoid the initialization process that occurs when you create a new one.

To create a consistency group snapshot virtual disk repository:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the consistency group to which you want to add member virtual disks and then select **Copy Services → Consistency Group → Remove Member Virtual Disks.**

   The **Consistency Group Snapshot Virtual Disk Settings** window is displayed.
3. Select **Manual** and click **Next** to customize the repository candidate settings for the consistency group.

   The **Consistency Group Snapshot Virtual Disk Repository Settings - Manual** window is displayed.
4. Select how you want filter the repository candidates for each member virtual disk in the consistency group, based on either a percentage of the base virtual disk capacity or by preferred capacity.

   The best repository candidate for each member virtual disk based on your selections is displayed.
5. Select **Edit individual repository candidates** if you want to edit repository candidates for the member virtual disks.
6. Select the repository, from the **Repository candidates** table, that you want to use for each member virtual disk in the consistency group.

   Select a repository candidate that is closest to the capacity you specified.

   – The **Repository candidates** table shows both new and existing repositories that are capable of being used for each member virtual disk in the consistency group based on the value you specified for percentage or the value you specified for preferred capacity.

   – By default, the system displays the repositories for each member virtual disk of the consistency group using a value of 20 percent of the member virtual disk's capacity. It filters out undersized repository candidates, and those with different Data Service (DS) attributes. If appropriate candidates are not returned using these settings, you can click **Run Auto-Choose** to provide automatic candidate recommendations.

– The **Difference** column shows the mathematical difference between your selected capacity and the actual capacity of the repository candidate. If the repository candidate is new, the system uses the exact capacity size that you specified and displays zero (0) in the **Difference** column.

7. To edit an individual repository candidate:
   a) Select the candidate from the **Repository candidates** table and click **Edit** to modify the capacity settings for the repository.
   b) Click **OK**.
8. In the **% full** box, define the value that determines when a warning is triggered when the capacity of a consistency group snapshot virtual disk repository reaches the defined percentage.
9. Click **Finish** to create the repository.

Disabling A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Disable** option when you want to invalidate a snapshot copy or a consistency group snapshot virtual disk. If the snapshot virtual disk or consistency group snapshot virtual disk is designated as read-write, this option also allows you stop any further write activity to its associated snapshot repository virtual disk.

Use the **Disable** option if one of these conditions applies:

- You are finished with the snapshot virtual disk or consistency group snapshot virtual disk for the time being.
- You intend to re-create the snapshot virtual disk or consistency group snapshot virtual disk (that is designated as read-write) at a later time and want to retain the associated snapshot repository virtual disk so that it does not need to be created again.
- You want to maximize the storage array performance by stopping write activity to the snapshot repository virtual disk.

If you decide to re-create the snapshot virtual disk or consistency group snapshot virtual disk, you must choose a snapshot image from the same base virtual disk.

If you disable the snapshot virtual disk or consistency group snapshot virtual disk, the system performs the following actions:

- Retains the World-Wide Name (WWN) for the snapshot virtual disk or consistency group snapshot virtual disk.
- Retains the snapshot virtual disk or consistency group snapshot virtual disk's association with the same base virtual disk.
- Retains the snapshot virtual disk or consistency group snapshot virtual disk's associated repository (if the virtual disk is designated as read-write).
- Retains any host mapping and access (any read-write requests will fail).
- Removes the snapshot virtual disk or consistency group snapshot virtual disk's association with the current snapshot image.
- For a consistency group snapshot virtual disk, disables each member's snapshot virtual disk.

NOTE: If you are finished with the snapshot virtual disk or consistency group snapshot virtual disk and do not intend to re-create it at a later time, you must delete the virtual disk, instead of disabling it.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:
   – **Copy Services → Snapshot Virtual disk → Disable.** The **Confirm Disable Snapshot Virtual Disk** window is displayed.
   – **Copy Services → Consistency Group Snapshot Virtual Disk → Disable.** The **Confirm Disable Consistency Group Snapshot Virtual Disk** window is displayed.
3. Type `yes` in the text box and then click **Disable** to disable the snapshot virtual disk.

The snapshot virtual disk or consistency group snapshot virtual disk is displayed in the Logical pane with the Disabled Snapshot status icon. If you disabled a read-write snapshot virtual disk or consistency group snapshot virtual disk, its associated snapshot repository virtual disk does not change status. The write activity to the snapshot repository virtual disk stops until the snapshot virtual disk or consistency group snapshot virtual disk is re-created.

Re-creating A Snapshot Virtual Disk Or Consistency Group Snapshot Virtual Disk

Use the **Re-Create** option when you want to re-create a snapshot virtual disk or consistency group snapshot virtual disk that you previously disabled. Re-creating a snapshot virtual disk or consistency group snapshot virtual disk takes less time than creating a new one.

If you have a snapshot virtual disk or consistency group snapshot virtual disk that you no longer need, you can reuse it (and any associated snapshot repository virtual disk), instead of deleting it, to create a different snapshot virtual disk or consistency group snapshot virtual disk of the same base virtual disk. You can re-associate the snapshot virtual disk or consistency group snapshot virtual disk with the same snapshot image or a different snapshot image as long as the snapshot image is in the same base virtual disk.

NOTE: If the snapshot virtual disk or consistency group snapshot virtual disk is part of an online copy relationship, you cannot perform the **Re-create** option on the virtual disk.

Keep these important guidelines in mind when you re-create a snapshot virtual disk or consistency group virtual disk:

- The snapshot virtual disk or consistency group snapshot virtual disk must be in either an **Optimal** status or **Disabled** status.
- For consistency group snapshot virtual disk, all member snapshot virtual disks must be in a Disabled state before you can re-create the consistency group snapshot virtual disk.
- You cannot re-create an individual member snapshot virtual disk, you can re-create only the overall consistency group snapshot virtual disk.
- All write data on any associated snapshot repository virtual disk is deleted. Snapshot virtual disk or consistency group snapshot virtual disk parameters remain the same as the previously disabled virtual disk parameters. The original names for the snapshot virtual disk or consistency group snapshot virtual disk are retained. You can change these names after the re-create option completes.

To re-create a snapshot virtual disk or consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot virtual disk or consistency group snapshot virtual disk that you want to disable and then select one of the following:
   - **Copy Services → Snapshot Virtual disk → Re-create.** The **Confirm Re-Create Snapshot Virtual Disk** window is displayed.
   - **Copy Services → Consistency Group Snapshot Virtual Disk → Re-create.** The **Confirm Re-Create Consistency Group Snapshot Virtual Disk** window is displayed.
3. Select whether to re-create the snapshot virtual disk or consistency group snapshot virtual disk using an existing snapshot image, or a new snapshot image and then click **Re-create**.

   The status of the snapshot virtual disk or consistency group snapshot virtual disk is changed from **Disabled** to **Optimal**.

Changing The Modification Priority Of An Overall Repository Virtual Disk

Use the **Modification Priority** option to specify the modification priority setting for an overall repository virtual disk on a storage array.

You can change the modification priority for an overall repository for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Replicated Pair

NOTE: Changing the modification priority by using this option modifies the priority only for the overall repository that you selected. The settings are applied to all individual repository virtual disks contained within the overall repository.

To change the modification priority:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the storage object for which to change the modification priority.
3. From the menu bar, select **Overall Repository → Change Modification Priority.**
   The **Change Disk Pool Settings** window is displayed.
4. In the **Select modification priority** area, move the slider bar to select a priority level.
5. Click **OK**.

Changing The Media Scan Setting Of An Overall Repository Virtual Disk

Use the **Change Media Scan Settings** option to set the media scan settings for an overall repository virtual disk on a storage array.

You can change the media scan settings for an overall repository for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Replicated pair

The following guidelines apply:

- Changing the media scan settings by using this option modifies the settings only for the overall repository that you selected.
- The settings are applied to all individual repository virtual disks contained within the overall repository.

To change the media scan settings:

1. In the AMW, select the **Storage & Copy Services** tab and select any virtual disk.
2. Select the storage object for which to change the media scan settings.
3. Select **Overall Repository → Change Media Scan Settings.**
   The **Change Media Scan Settings** window is displayed.
4. Select **Enable media scan**.
5. Select either **With consistency check** or **Without consistency check**, and click **OK**.
   A consistency check scans the blocks in a RAID Level 5 virtual disk, or a RAID Level 6 virtual disk and checks the consistency information for each block. A consistency check compares data blocks on RAID Level 1 mirrored physical disks. RAID Level 0 virtual disks have no data redundancy.

Changing The Pre-read Consistency Check Setting Of An Overall Repository Virtual Disk

Use the **Pre-Read Consistency Check** option to define a storage array's capability to pre-read an overall repository virtual disk consistency information and determine whether the data of that overall repository virtual disk is consistent.

An overall repository virtual disk that has this feature enabled returns read errors if the data is determined to be inconsistent by the RAID controller module firmware. You can enable this option for overall repository virtual disks that contain consistency information. RAID Level 1, RAID Level 5, and RAID Level 6 maintain redundancy information.

You can enable this option for overall repository virtual disks that contain consistency information. RAID Level 1, RAID Level 5, and RAID Level 6 maintain redundancy information.

You can change the Pre-Read Consistency Check for an overall repository for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Replicated Pair

The following guidelines apply:

- Changing the **Pre-Read Consistency Check** setting modifies the setting only for the overall repository that you selected.
- The **Pre-Read Consistency Check** setting is applied to all individual repository virtual disks contained within the overall repository.
- If an overall repository virtual disk that is configured with pre-read is migrated to a RAID level that does not maintain consistency information, the metadata of the overall repository virtual disk continues to show that pre-read is enabled. However, reads to that overall repository virtual disk ignores consistency pre-read. If the virtual disk is subsequently migrated back to a RAID level that supports consistency, the option becomes available again.

To create a consistency group snapshot virtual disk:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select **Overall Repository → Change Pre-read Consistency Check.**
3. Select **Enable pre-read consistency check**, and click **OK**.

   NOTE: Enabling the option on overall repository virtual disks without consistency does not affect the virtual disk. However, the attribute is retained for that overall repository virtual disk if it is ever changed to one with consistency information.
4. Click **Yes**.

Increasing The Capacity Of An Overall Repository

An overall repository can contain multiple repository virtual disks. You can use the **Increase Capacity** option to increase the storage capacity of an existing overall repository for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Consistency group member snapshot virtual disk
- Replicated pair

Use this option when you receive a warning that the overall repository is in danger of becoming full. You can increase the repository capacity by performing one of these tasks:

- Adding one or more existing repository virtual disks.
- Creating a new repository virtual disk using free capacity that is available on a disk group or disk pool.

NOTE: If no free capacity exists on any disk group or disk pool, you can add unconfigured capacity in the form of unused physical disks to a disk group or disk pool.

You cannot increase the storage capacity of an overall repository if one of these conditions exists:

- The repository virtual disk that you want to add does not have an Optimal status.
- Any repository virtual disk in the disk group or disk pool that you want to add is in any state of modification.
- No free capacity exists in the disk group or disk pool that you want to add.
- No unconfigured capacity exists in the disk group or disk pool that you want to add.
- There are no eligible existing repository virtual disks (including mismatched DS attributes).
- Make sure that a base virtual disk and each of the individual repository virtual disks in the overall repository have the same Data Service (DS) attributes, specifically for the following characteristics:
- RAID Level — A repository in a disk pool is considered to have a matching RAID Level for any base virtual disk on a disk group, regardless of the base virtual disk's actual RAID Level. However, a repository on a disk group is considered to have a matching RAID Level only if that RAID Level is identical to the RAID Level of the base virtual disk.
- Physical Disk Type — A match requires that the base virtual disk and the repository virtual disk reside on either a disk group or disk pool with identical physical disk type attributes.
- You cannot increase or decrease the repository capacity for a snapshot virtual disk that is read-only because it does not have an associated repository. Only snapshot virtual disks that are read-write require a repository.

To increase the overall repository capacity:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the storage object for which you want to increase the repository capacity and then select **Overall Repository → Increase Capacity.**

   The **Increase Repository Capacity** window is displayed.
3. To increase capacity of the overall repository, do one of the following:
   - Select **Add one or more existing repository virtual disks** and then go to step 4.
   - Select **Create and add new repository virtual disk** and then go to step 5.
4. To add one or more existing repository virtual disks, perform the following steps:
   a) Select one or more repository virtual disks from the **Eligible repository virtual disks** table.

   The eligible repository virtual disks that have the same DS settings as the associated base virtual disk are only displayed.

   NOTE: You can click the **Select all** check box to add all the repository virtual disks displayed in the **Eligible repository virtual disks** table.

   b) Select **Allow mismatch in DS attributes** to display additional repository virtual disks that do not have the same DS settings as the base virtual disk.
5. To create a new repository virtual disk, perform the following steps:
   a) From the **Create New Repository On** drop-down list, select a disk group or disk pool.

   The drop-down lists only the eligible repository virtual disks that have the same DS settings as the associated base virtual disk. You can select **Allow mismatch in DS attributes** to display additional repository virtual disks that do not have the same DS settings as the base virtual disk.

   If free capacity is available in the disk group or disk pool you selected, the total free space is displayed in the **Capacity** spinner box.

   b) If required, adjust the **Capacity**.

NOTE: If free capacity does not exist on the disk group or disk pool you selected, the free space that appears in the **Capacity** spinner box is 0. If this storage array has **Unconfigured Capacity**, you can create a new disk group or disk pool and then retry this operation using the new free capacity on that disk group or disk pool.

6. Click **Increase Repository**.

   The system performs the following actions:

   - Updates the capacity for the repository.
   - Displays the newly-added repository member virtual disk(s) for the repository.

Decreasing The Capacity Of The Overall Repository

An overall repository can contain multiple repository virtual disks.

Use the **Decrease Capacity** option to decrease the storage capacity of an existing overall repository for the following storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Consistency group member snapshot virtual disk
- Replicated pair virtual disk

You cannot decrease the storage capacity of the overall repository if one of these conditions exists:

- The overall repository contains only one repository member virtual disk.
- If there are one or more snapshot images associated with the overall repository.
- If a snapshot virtual disk or a consistency group member snapshot virtual disk is disabled.

The following guidelines apply:

- You can remove repository member virtual disks only in the reverse order that they were added.
- An overall repository must have at least one repository member virtual disk.
- You cannot increase or decrease the repository capacity for a snapshot virtual disk that is read-only because it does not have an associated repository. Only snapshot virtual disks that are read-write require a repository.
- When you decrease capacity for a snapshot virtual disk or a consistency group member snapshot virtual disk, the system automatically transitions the virtual disk to a **Disabled** state.

To decrease the overall repository capacity:

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the storage object for which you want to decrease the repository capacity and then select **Overall Repository → Decrease Capacity.**

   The **Decrease Repository Capacity** window is displayed.
3. Select one or more repository virtual disks from the **Repository member virtual disks** table that you want to remove.

   - The table displays the member virtual disks in reverse order that they were added for the storage object. When you can click on any row in the table, that row and all rows above it are selected.
   - The last row of the table, which is the first repository added, is disabled because at least one repository must exist for the storage object.
4. Click **Delete selected repository virtual disks** if you want to delete all associated repositories that exist for each member virtual disk selected in the **Repository member virtual disks** table.
5. Click **Decrease Repository**.

The system performs the following actions:

- Updates the capacity for the overall repository.
- Displays the newly-updated repository member virtual disk(s) for the overall repository.

Performing A Revive Operation

Use the **Revive** option to force a storage object to an Optimal state if it does not transition automatically after a failure is corrected.

You can use the **Revive** option for these storage objects:

- Snapshot group
- Snapshot virtual disk
- Consistency group member virtual disk
- Consistency group member snapshot virtual disk

NOTE: Use the **Revive** option only if you are instructed to do so in a Recovery Guru procedure or by a Technical Support representative. You cannot cancel this operation after it starts.

Use this option when you receive a warning that the overall repository is in danger of becoming full. You can increase the repository capacity by performing one of these tasks:

- Adding one or more existing repository virtual disks.
- Creating a new repository virtual disk using free capacity that is available on a disk group or disk pool.

  NOTE: If no free capacity exists on any disk group or disk pool, you can add unconfigured capacity in the form of unused physical disks to a disk group or disk pool.

  CAUTION: Using the Revive option when there are still failures may cause data corruption or data loss, and the storage object will return to the Failed state.

1. From the AMW, select the **Storage & Copy Services** tab.
2. Select the storage object that you want to revive and then select one of the following menu paths (depending on the storage object you selected):
   - **Copy Services → Snapshot Group → Advanced → Revive.**
   - **Copy Services → Snapshot Virtual Disk → Advanced → Revive.**
   - **Copy Services → Consistency Group Member Virtual Disk → Advanced → Revive.**
3. Type `yes` in the text box and then click **Revive** to restore the storage object to an **Optimal** state.

11

Premium Feature—Snapshot Virtual Disks (Legacy)

The following types of virtual disk snapshot premium features are supported on the MD storage array:

- Snapshot Virtual Disks using multiple point-in-time (PiT) groups
- Snapshot Virtual Disks (Legacy) using a separate repository for each snapshot

NOTE: This section describes the Snapshot Virtual Disk (Legacy) premium feature.

NOTE: If you ordered this feature, you received a Premium Feature Activation card shipped in the same box as your Dell PowerVault MD storage array. Follow the directions on the card to obtain a key file and to enable the feature.

NOTE: The snapshot feature allows up to eight snapshots per LUN and 128 per array to be present at the same time.

A snapshot virtual disk is a point-in-time image of a virtual disk in a storage array. It is not an actual virtual disk containing a copy of the original data; It is a reference to the data that was contained on a virtual disk at a specific time. A snapshot virtual disk is the logical equivalent of a complete physical copy. However, you can create a snapshot virtual disk much faster than a physical copy, using less disk space.

The virtual disk on which the snapshot is based, called the source virtual disk, must be a standard virtual disk in your storage array. Typically, you create a snapshot so that an application, such as a backup application, can access the snapshot and read the data while the source virtual disk remains online and accessible.

NOTE: No I/O requests are permitted on the source virtual disk while the virtual disk snapshot is being created.

A snapshot repository virtual disk containing metadata and copy-on-write data is automatically created when a snapshot virtual disk is created. The only data stored in the snapshot repository virtual disk is that which has changed since the time of the snapshot.

After the snapshot repository virtual disk is created, I/O write requests to the source virtual disk resume. Before a data block on the source virtual disk is modified, the contents of the block to be modified are copied to the snapshot repository virtual disk for safekeeping. Because the snapshot repository virtual disk stores copies of the original data in those data blocks, further changes to those data blocks write only to the source virtual disk. The snapshot repository uses less disk space than a full physical copy, because the only data blocks that are stored in the snapshot repository virtual disk are those that have changed since the time of the snapshot.

When you create a snapshot virtual disk, you specify its location, capacity, and other parameters. You can disable or delete the snapshot virtual disk when it is not required. If you disable a snapshot virtual disk, you can re-create and reuse it the next time you perform a backup. If you delete a snapshot virtual disk, you also delete the associated snapshot repository virtual disk.

NOTE: If the Source Virtual Disk is in the offline state, the corresponding Snapshot(s) Repository(ies) and Snapshot(s) Virtual Disk(s) will be in **Failed** state.

NOTE: Deleting a snapshot does not affect data on the source virtual disk.

NOTE: The following host preparation sections also apply when using the snapshot feature through the CLI interface.

Scheduling A Snapshot Virtual Disk

When you create a snapshot virtual disk, you can choose whether the snapshot is created immediately or is created according to a schedule that you determine. This schedule can be a one-time snapshot creation or an snapshot creation that occurs at regularly occurring intervals. If a schedule is not specified, the snapshot virtual disk creation happens immediately upon execution of the command.

A schedule can be specified when a snapshot virtual disk is first created, or it can be added to an existing snapshot virtual disk at any time. One schedule per snapshot virtual disk is supported.

Common Reasons For Scheduling A Snapshot Virtual Disk

Scheduling a snapshot virtual disk can serve multiple purposes across a data storage environment. Most common uses of a snapshot scheduler are:

- Data backups
- Rapid recovery from a data loss event

A scheduled data backup can protect against data loss on a regular, unmonitored basis. For example, if an application stores business-critical data on two virtual disks in the storage array, you may choose to perform an automatic backup every day. To implement this backup, select the first virtual disk and create a backup schedule that runs once a day, Monday through Friday, at a time between the end of the work day and 11PM. Do not select an end date. Apply the same schedule to the second virtual disk, then map the two snapshot virtual disks to your backup host server and perform your regular backup procedures. Remember to unmap the two resulting snapshot virtual disks before the next scheduled snapshot begins. If the snapshot virtual disks are not unmapped, the storage array does not perform the next scheduled snapshot operation in order to avoid data corruption.

Scheduled snapshots are also valuable in the event of a data loss. For example, if you back up your data at the end of every work day and keep hourly snapshots from 8AM to 5PM, data can be recovered from the snapshots in windows smaller than one hour. To accomplish this type of rapid recovery, create a schedule that contains a start time of 8AM and an end time of 5PM, then select 10 snapshots per day on Monday through Friday with no end date.

Guidelines for Creating Snapshot Schedules

Certain guidelines apply when creating snapshot virtual disk schedules:

- Scheduled virtual disk snapshot operations do not occur if:
  - The snapshot virtual disk is mapped
  - The storage array is offline or turned off
  - The snapshot virtual disk is in use as a source virtual disk during a Virtual Disk Copy operation
  - A copy operation is Pending or In progress
- Deleting a snapshot virtual disk that contains a schedule also deletes the schedule.
- Snapshot schedules are stored in the configuration database on the storage array. The Management Station does not need to be running for scheduled snapshot operations to occur.
- Snapshot schedules can be created when the snapshot virtual disk is initially created or can be added to existing snapshot virtual disks.

Creating A Snapshot Virtual Disk Using The Simple Path

You can choose the simple path to create a snapshot virtual disk if the disk group of the source virtual disk has the required amount of free space. A snapshot repository virtual disk requires a minimum of 8 MB free capacity. The

destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group. If 8 MB of free capacity is not available in the disk group of the source virtual disk, the Create Snapshot Virtual Disks feature defaults to the advanced path. In the advanced path option, you can choose to place the snapshot repository virtual disk in another disk group or you can use unconfigured capacity on the storage array to create a new disk group. For more information, see Creating A Snapshot Virtual Disk Using The Advanced Path.

About The Simple Path

Using the simple path, you can specify:

- Snapshot Virtual Disk Name — A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.
- Snapshot Repository Virtual Disk Name — A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.
- Snapshot Repository Virtual Disk Capacity — The snapshot repository virtual disk capacity is expressed as a percentage of the source virtual disk capacity. The maximum percentage allowed is 120 percent.

Using the simple path, the following defaults are used for the other parameters of a snapshot virtual disk:

- Capacity Allocation — The snapshot repository virtual disk is created using free capacity on the same disk group where the source virtual disk resides.
- Host-to-Virtual Disk Mapping — The default setting is **Map now**.
- Percent Full — When the snapshot repository virtual disk reaches the specified repository full percentage level, the event is logged in the Major Event Log (MEL). The default snapshot repository full percentage level is 50 percent of the source virtual disk.
- Snapshot Repository Virtual Disk Full Conditions — When the snapshot repository virtual disk becomes full, you are given a choice of failing write activity to the source virtual disk or failing the snapshot virtual disk.

Preparing Host Servers To Create The Snapshot Using The Simple Path

- Before using the Snapshot Virtual Disks Premium Feature in a Microsoft Windows clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.
- Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.
- For details on mapping the snapshot virtual disk to the secondary node, see the documentation on storage arrays with Microsoft Windows Server Failover Clusters at **dell.com/support/manuals**.
- You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Before creating a Snapshot Virtual Disk, note that:

- The following types of virtual disks are not valid source virtual disks:
  - Snapshot repository virtual disks
  - Snapshot virtual disks
  - Target virtual disks that are participating in a virtual disk copy
- You cannot create a snapshot of a virtual disk that contains unreadable sectors.
- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate snapshot of the source virtual disk or the target virtual disk in a virtual disk copy.

NOTE: Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before creating a snapshot virtual disk, the host server has to be in the proper state. To ensure that the host server is properly prepared to create a snapshot virtual disk, you can either use an application to carry out this task, or you can perform the following steps:

1. Stop all I/O activity to the source.
2. In the AMW, select the **Storage & Copy Services** tab and select a valid source virtual disk.
3. Select **Copy Services → Snapshot Image → Create Snapshot Virtual Disk.** Alternatively, you can right-click the source virtual disk and select **Create → Snapshot Virtual Disk** from the pop-up menu.

   The **Create Snapshot Virtual Disk Wizard - Introduction** dialog is displayed.
4. Select **Simple (Recommended)**, and click **Next**.

   The **Specify Names** window is displayed.
5. Enter the **Snapshot virtual disk name** and the **Snapshot repository virtual disk name** and click **Next**.

   The **Specify Snapshot Repository Capacity** window is displayed.
6. Enter the snapshot repository virtual disks capacity as a percentage of the source virtual disks capacity and click **Next**.

   The **Preview** window containing the summary of the snapshot virtual disk is displayed.
7. Click **Finish**.

   The **Completed** window is displayed.
8. Click **OK**.

   After creating one or more snapshot virtual disks, mount the source virtual disk, and restart the host application using that source virtual disk.
9. In the AMW, select the **Host Mappings** tab, assign mappings between the snapshot virtual disk and the host that accesses the snapshot virtual disk.

   NOTE: In some cases, conflicts might result from mapping the same host to both a source virtual disk and its associated snapshot virtual disk. This conflict depends on the host operating system and any virtual disk manager software in use.
10. To register the snapshot virtual disk with the host operating system, run the host-based hot_add utility.
11. To associate the mapping between the storage array name and the virtual disk name, run the host-based SMdevices utility

    NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Creating A Snapshot Virtual Disk Using The Advanced Path

About The Advanced Path

Use the advanced path to choose whether to place the snapshot repository virtual disk on free capacity or unconfigured capacity and to change the snapshot repository virtual disk parameters. You can select the advanced path regardless of whether you use free capacity or unconfigured capacity for the snapshot virtual disk.

Using the advanced path, you can specify the following parameters for your snapshot virtual disk:

- Snapshot Virtual Disk Name — A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.
- Snapshot Repository Virtual Disk Name — A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.
- Capacity Allocation — This parameter allows you to choose where to create the snapshot repository virtual disk. You can allocate capacity by using one of the following methods:
  - Use free capacity on the same disk group where the source virtual disk resides.
  - Use free capacity on another disk group.
  - Use unconfigured capacity and create a new disk group for the snapshot repository virtual disk.
  - It is recommended placing the snapshot repository virtual disk within the disk group of the source virtual disk. This ensures that if drives associated with the disk group are moved to another storage array, all the virtual disks associated with the snapshot virtual disk remain in the same group.
- Snapshot Repository Virtual Disk Capacity — The snapshot repository virtual disk capacity is expressed as a percentage of the source virtual disk capacity. The maximum percentage allowed is 120 percent.
- Percent Full—When the snapshot repository virtual disk reaches the user-specified repository full percentage level, the event is logged in the Major Event Log (MEL). The default snapshot repository full percentage level is 50% of the source virtual disk.
- Snapshot Repository Virtual Disk Full Conditions — Choose whether to fail writes to the source virtual disk or fail the snapshot virtual disk when the snapshot repository virtual disk becomes full.
- Host-to-Virtual Disk Mapping — Choose whether to map the snapshot virtual disk to a host or host group now or to map the snapshot virtual disk later. The default setting is **Map later**.

Preparing Host Servers To Create The Snapshot Using The Advanced Path

- Before using the Snapshot Virtual Disks Premium Feature in a Microsoft Windows clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.
- Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.
- For details on mapping the snapshot virtual disk to the secondary node, see the documentation on storage arrays with Microsoft Windows Server Failover Clusters at **dell.com/support/manuals**.

The destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group. A snapshot repository virtual disk requires a minimum of 8 MB free capacity. You can choose your preferred creation path—simple or advanced—if the disk group of the source virtual disk has the required amount of free space.

If 8 MB of free capacity is not available in the disk group of the source virtual disk, the Create Snapshot Virtual Disks feature defaults to the advanced path. See Creating A Snapshot Virtual Disk Using The Advanced Path.

NOTE: You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Before creating a Snapshot Virtual Disk, note that:

- The following types of virtual disks are not valid source virtual disks:
  - Snapshot repository virtual disks
  - Snapshot virtual disks
  - Target virtual disks that are participating in a virtual disk copy
- You cannot create a snapshot of a virtual disk that contains unreadable sectors.
- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate snapshot of the source virtual disk or the target virtual disk in a virtual disk copy.

NOTE: Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before creating a snapshot virtual disk, the host server must be in the proper state. To prepare your host server:

1. Stop all I/O activity to the source.
2. Using your Windows system, flush the cache to the source. At the host prompt, type `SMrepassist -f <filename-identifier>` and press <Enter>. For more information, see SMrepassist Utility.
3. Remove the drive letter(s) of the source in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot.

   If this is not done, the snapshot operation reports that it has completed successfully, but the snapshot data is not updated properly.

   NOTE: To verify that the virtual disk is in Optimal or Disabled state, select the **Summary** tab and then click **Disk Groups & Virtual Disks**.

   The **Create Snapshot Virtual Disk Wizard - Introduction** dialog is displayed.
4. Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

   NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

   After your host server has been prepared, see Creating A Snapshot Virtual Disk Using The Advanced Path, to create the snapshot using the advanced path.

   If you want to use a snapshot regularly, such as for backups, use the Disable Snapshot and Re-create Snapshot options to reuse the snapshot. Disabling and re-creating snapshots preserves the existing virtual disk-to-host mappings to the snapshot virtual disk.

Creating The Snapshot Using The Advanced Path

NOTE: Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Prepare the host server(s) as specified in Preparing Host Servers To Create The Snapshot Using The Advanced Path.

To create a virtual disk snapshot using the advanced path:

1. Stop the host application accessing the source virtual disk, and unmount the source virtual disk.
2. In the AMW, select the **Storage & Copy Services** tab, select a valid source virtual disk.
3. Select **Copy Services → Snapshot Image → Create Snapshot Virtual Disk.** Alternatively, right-click the source virtual disk and select **Create → Snapshot Virtual Disk** from the pop-up menu.

   The **Create Snapshot Virtual Disk Wizard - Introduction** dialog is displayed.
4. Select **Advanced**, and click **Next**.

   The **Specify Names** window is displayed.
5. Enter the **Snapshot virtual disk name** and the **Snapshot repository virtual disk name** and click **Next**.

   The **Allocate Capacity** window is displayed.
6. In the **Capacity allocation** area, select:

- Free capacity on same disk group as base (recommended)
- Free capacity on different disk group
- Unconfigured capacity (create new disk group)

7. Enter the snapshot repository virtual disks capacity as a percentage of the source virtual disks capacity and click **Next**.

   The **Specify Virtual Disk Parameters** window is displayed.

8. In the **Snapshot virtual disk parameters** area, select the relevant mapping option, you can select:

   - **Automatic**
   - **Map later with Storage Partition**

9. In the **Snapshot repository virtual disk parameters** area, enter the system behavior when:
10. Click **Next**.

    The **Preview** window containing the summary of the snapshot virtual disk is displayed.

11. Click **Finish**.

    The **Completed** window is displayed.

12. Click **OK**.
13. In the **Host Mappings** tab, assign mappings between the snapshot virtual disk and the host that accesses the snapshot virtual disk.
14. To register the snapshot virtual disk with the host operating system, run the host-based hot_add utility.
15. To associate the mapping between the storage array name and the virtual disk name, run the host-based SMdevices utility.

Specifying Snapshot Virtual Disk Names

Choose a name that helps you associate the snapshot virtual disk and snapshot repository virtual disk with its corresponding source virtual disk. The following information is useful when naming virtual disks.

By default, the snapshot name is shown in the Snapshot virtual disk name field as:

```
<source-virtual disk-name>—<sequence-number>
```

where `sequence-number` is the chronological number of the snapshot relative to the source virtual disk.

The default name for the associated snapshot repository virtual disk that is shown in the Snapshot repository virtual disk field is:

```
<source-virtual disk-name>—R<sequence-number>
```

For example, if you are creating the first snapshot virtual disk for a source virtual disk called Accounting, the default snapshot virtual disk is Accounting-1, and the associated snapshot repository virtual disk default name is Accounting-R1. The default name of the next snapshot virtual disk you create based on Accounting is Accounting-2, with the corresponding snapshot repository virtual disk named as Accounting-R2 by default.

- Whether you use the software-supplied sequence number that (by default) populates the Snapshot virtual disk name or the Snapshot repository virtual disk name field, the next default name for a snapshot or snapshot repository virtual disk still uses the sequence number determined by the software. For example, if you give the first snapshot of source virtual disk Accounting the name Accounting-8, and do not use the softwaresupplied sequence number of 1, the default name for the next snapshot of Accounting is still Accounting-2.
- The next available sequence number is based on the number of existing snapshots of a source virtual disk. If you delete a snapshot virtual disk, its sequence number becomes available again.
- You must choose a unique name for the snapshot virtual disk and the snapshot repository virtual disks, or an error message is displayed.
- Names are limited to 30 characters. After you reach this limit in either the snapshot virtual disk name or the Snapshot repository virtual disk name fields, you can no longer type in the field. If the source virtual disk is 30

characters, the default names for the snapshot and its associated snapshot repository virtual disk use the source virtual disk name truncated enough to add the sequence string. For example, for Host Software Engineering Group GR-1, the default snapshot name is Host Software Engineering GR- 1, and the default repository name is Host Software Engineering GR-R1.

Snapshot Repository Capacity

If you receive a warning that the capacity for the snapshot repository virtual disk is approaching its threshold, you can increase the capacity of a snapshot repository virtual disk by using one of the following methods:

- Use the free capacity available on the disk group of the snapshot repository virtual disk.
- Add unconfigured capacity to the disk group of the snapshot repository virtual disk. Use this option when no free capacity exists on the disk group.

You cannot increase the storage capacity of a snapshot repository virtual disk if the snapshot repository virtual disk has any one of the following conditions:

- The virtual disk has one or more hot spare drives in use.
- The virtual disk has a status other than Optimal.
- Any virtual disk in the disk group is in any state of modification.
- The controller that has ownership of this virtual disk is currently adding capacity to another virtual disk. Each controller can add capacity to only one virtual disk at a time.
- No free capacity exists in the disk group.
- No unconfigured capacity is available to add to the disk group.

  NOTE: You can add a maximum of two physical disks at one time to increase snapshot repository virtual disk capacity.

To expand the snapshot repository virtual disk from MD Storage Manager:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select the snapshot repository virtual disk for which you want to increase the capacity.
3. From the menu bar, select **Storage → Virtual Disk → Increase Capacity.** Alternatively, right-click on the snapshot repository virtual disk and select **Increase Capacity**.

   NOTE: If no free capacity or unconfigured capacity is available, the **Increase Capacity** option is disabled.

   The **Increase Snapshot Repository Capacity** window displays the Virtual disk attributes. The snapshot repository virtual disk name, the associated snapshot virtual disk name, the associated source virtual disk capacity and name, the current capacity, and the amount of free capacity that is available for the selected snapshot repository virtual disk are displayed. If free capacity is available, the maximum free space is displayed in **Increase capacity by**.

   If free capacity does not exist on the disk group, the free space that is displayed in **Increase capacity by** is 0. You must add physical disks to create free capacity on the disk group.
4. To increase capacity of the snapshot repository virtual disk, use one of these methods:
   - Use the free capacity on the disk group of the snapshot repository virtual disk—Go to step 5.
   - Add unconfigured capacity, or physical disks to the disk group of the snapshot repository virtual disk—Go to step 7.
5. In **Increase capacity by**, enter or select the appropriate capacity.
6. Click **OK**.

   The **Storage & Copy Services** tab is updated. The snapshot repository virtual disk having its capacity increased shows a status of Operation in Progress. In addition, the snapshot repository virtual disk shows its original capacity and the total capacity being added. The virtual disk involved shows a reduction in capacity. If all of the free

capacity is used to increase the size of the virtual disk, the Free Capacity node involved is removed from the **Storage & Copy Services** tab.

7. If unassigned physical disks are not available, do you have empty slots in the expansion enclosures?
   - Yes, there are empty slots—Insert new physical disks by using the information in the initial setup guide for your expansion enclosure. Go to step 9.
   - No, there are no empty slots—Install another expansion enclosure and additional physical disks. Use the information in the initial setup guides for your RAID controller module and your expansion enclosure. Go to step 9.

     NOTE: The physical disks that you add must be of the same media type and interface type as the physical disks that already make up the disk group of the snapshot repository virtual disk.
8. Click **Add Physical Disks**.
9. Select either a single physical disk to add or two physical disks to add.
10. Click **Add**.

    The **Add Physical Disks** window closes.
11. Check the **Physical Disks to add** [enclosure, slot] area to make sure that the correct physical disks have been added.
12. Either accept the final capacity, or enter or select the appropriate capacity in **Increase capacity by**.
13. Click **OK**.

    The **Storage & Copy Services** tab is updated. The snapshot repository virtual disk that is having its capacity increased shows a status of Operation in Progress. In addition, the snapshot repository virtual disk shows its original capacity and the total capacity being added. The Free Capacity node involved in the increase shows a reduction in capacity. If all of the free capacity is used to increase the size of the virtual disk, the Free Capacity node involved is removed from the **Storage & Copy Services** tab.

    A new Free Capacity node is created and shown in the **Storage & Copy Services** tab if these conditions exist:
    - A Free Capacity node did not exist prior to the addition of capacity.
    - Not all of the capacity that is added is used to increase the capacity of the snapshot repository virtual disk.

    On the **Hardware** tab, the unassigned physical disks or unconfigured capacity that you added to increase the capacity of the snapshot repository virtual disk change to assigned physical disks. The new assigned physical disks are associated with the disk group of the snapshot repository virtual disk.

Re-Creating Snapshot Virtual Disks

You can re-create a snapshot virtual disk that you have previously disabled.

CAUTION: Possible loss of data redundancy—If the snapshot virtual disk is in Optimal status, it is first disabled prior to being re-created. This action invalidates the current snapshot.

Keep these important guidelines in mind when you re-create a snapshot virtual disk:

- To re-create the snapshot virtual disks correctly, follow the instructions for your operating system.

  NOTE: Failing to follow these additional instructions could create unusable snapshot virtual disks. For more information, see the online help topics.
- To use this option, the snapshot virtual disk must be either in an Optimal status or Disabled status.
- When using this option, the previously configured snapshot name parameters and snapshot repository virtual disk are used.

To recreate the snapshot virtual disk:

1. In the AMW, select the **Storage & Copy Services** tab, select a snapshot virtual disk.
2. Select **Copy services → Virtual Disk → Snapshot → Re-create.**
3. Type `yes`, and click **OK**.

Disabling A Snapshot Virtual Disk

Disable a snapshot virtual disk if one of the following conditions exists:

- You do not need the snapshot now.
- You intend to re-create the snapshot at a later time and want to retain the associated snapshot repository virtual disk so that you do not need to create it again.
- You want to maximize storage array performance by stopping copy-onwrite activity to the snapshot repository virtual disk.

NOTE: If you do not intend to re-create the snapshot virtual disk at a later time, in the object tree, select the snapshot virtual disk, and select **Virtual Disk → Delete** to remove it. The associated snapshot repository virtual disk is also removed. See the online help topics for more information on removing a snapshot virtual disk.

NOTE: The SMdevices utility displays the snapshot virtual disk in its output, even after the snapshot virtual disk is disabled.

To disable a snapshot virtual disk:

1. In the AMW, select the **Storage & Copy Services** tab, select the snapshot virtual disk, and select **Copy Services → Snapshot Virtual Disk → Disable.**
2. In the text box, type `yes` and click **OK**.

   The snapshot virtual disk is disabled. The associated snapshot repository virtual disk does not change status. The copy-on-write activity to the snapshot repository virtual disk stops until the snapshot virtual disk is recreated.

Preparing Host Servers To Re-Create A Snapshot Virtual Disk

NOTE: Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk and snapshot virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before re-creating a snapshot virtual disk, both the host server and the associated virtual disk you are re-creating have to be in the proper state.

To prepare your host server and virtual disk:

1. Stop all I/O activity to the source and snapshot virtual disk (if mounted).
2. Using your Windows system, flush the cache to both the source and the snapshot virtual disk (if mounted). At the host prompt, type `SMrepassist -f <filename-identifier>` and press <Enter>.

   For more information, see SMrepassist Utility.
3. Click the **Summary** tab, then click **View Storage Array Profile** to verify if the snapshot virtual disk is in Optimal or Disabled status.
4. Remove the drive letter(s) of the source and (if mounted) snapshot virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation reports that it has completed successfully, but the snapshot data is not updated properly.

5. Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

   NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Re-Creating A Snapshot Virtual Disk

After first preparing the host server(s), re-create a snapshot virtual disk. For more information, see Preparing Host Servers To Create The Snapshot Using The Simple Path or Preparing Host Servers To Create The Snapshot Using The Advanced Path.

NOTE: This action invalidates the current snapshot.

To recreate a snapshot virtual disk:

1. In the AMW, select the **Storage & Copy Services** tab and select a snapshot virtual disk.
2. Select **Copy Services → Snapshot Virtual Disk → Re-Create.**
3. Type `yes`, and click **OK**.

   Re-creating a snapshot repository virtual disk uses the previously configured snapshot name and parameters.

12

Premium Feature—Virtual Disk Copy

NOTE: A virtual disk copy overwrites data on the target virtual disk. Before starting a virtual disk copy, ensure that you no longer need the data or back up the data on the target virtual disk.

NOTE: If you ordered this feature, you received a Premium Feature Activation card that shipped in the same box as your Dell PowerVault MD Series storage array. Follow the directions on the card to obtain a key file and to enable the feature.

NOTE: The preferred method for creating a virtual disk copy is to copy from a snapshot virtual disk. This allows the original virtual disk used in the snapshot operation to remain fully available for read/write activity while the snapshot is used as the source for the virtual disk copy operation.

When you create a virtual disk copy, you create a copy pair that has a source virtual disk and a target virtual disk on the same storage array.

The source virtual disk is the virtual disk that contains the data you want to copy. The source virtual disk accepts the host I/O read activity and stores the data until it is copied to the target virtual disk. The source virtual disk can be a standard or thin virtual disk.

The target virtual disk is a standard or thin virtual disk in a disk group or disk pool and, if the legacy version is enabled, a legacy snapshot base virtual disk.

Reasons to use virtual disk copy include:

- Copying data for improved access — As your storage requirements for a virtual disk change, you can use a virtual disk copy to copy data to a virtual disk in a disk group that uses drives with larger capacity within the same storage array. Copying data for larger access capacity enables you to move data to greater capacity physical disks (for example, 61 GB to 146 GB).
- Restoring snapshot virtual disk data to the source virtual disk — The Virtual Disk Copy feature enables you first to restore the data from a snapshot virtual disk and then to copy the data from the snapshot virtual disk to the original source virtual disk.
- Copying data from a thin virtual disk to a standard virtual disk residing in the same storage array. However, you cannot copy data in the opposite direction (from a standard virtual disk to a thin virtual disk).
- Creating a backup copy — The Virtual Disk Copy feature enables you to create a backup of a virtual disk by copying data from one virtual disk (the source virtual disk) to another virtual disk (the target virtual disk) in the same storage array, minimizing the time that the source virtual disk is unavailable to host write activity. You can then use the target virtual disk as a backup for the source virtual disk, as a resource for system testing, or to copy data to another device, such as a tape drive or other media.

  NOTE: Recovering from a backup copy—You can use the Edit Host-to-Virtual Disk Mappings feature to recover data from the backup virtual disk you created in the previous procedure. The **Host Mappings** option enables you to unmap the source virtual disk from its host and then to map the backup virtual disk to the same host.

Using Virtual Disk Copy With Snapshot Or Snapshot (Legacy) Premium Feature

After completion of the virtual disk copy of a snapshot (Legacy), the legacy snapshot is disabled. After completion of the virtual disk copy using a snapshot image, the snapshot image is deleted and the snapshot virtual disk is disabled.

Snapshots created using older (Legacy) premium feature versions cannot be managed using newer snapshot premium feature options. Also, a virtual disk in a snapshot group cannot be a target for a virtual disk copy. If you want to choose

the base virtual disk of an older (Legacy) snapshot virtual disk as your target virtual disk, you must first disable all snapshot (Legacy) virtual disks that are associated with the base virtual disk.

Types Of Virtual Disk Copies

You can perform either offline or online virtual disk copies. To ensure data integrity, all I/O to the target virtual disk is suspended during either type of virtual disk copy operation. After the virtual disk copy is complete, the target virtual disk automatically becomes read-only to the hosts.

Offline Copy

An offline copy reads data from the source virtual disk and copies it to a target virtual disk, while suspending all updates to the source virtual disk when the copy is in progress. In an offline virtual disk copy, the relationship is between a source virtual disk and a target virtual disk. Source virtual disks that are participating in an offline copy are available for read requests, while the virtual disk copy displays the **In Progress** or **Pending** status. Write requests are allowed only after the offline copy is complete. If the source virtual disk is formatted with a journaling file system, any attempt to issue a read request to the source virtual disk may be rejected by the storage array RAID controller modules and result in an error message. Make sure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy is complete to prevent error messages from being displayed.

Online Copy

An online copy creates a point-in-time snapshot copy of any virtual disk within a storage array, while still allowing writes to the virtual disk when the copy is in progress. This is achieved by creating a snapshot of the virtual disk and using that snapshot as the actual source virtual disk for the copy. In an online virtual disk copy, the relationship is between a snapshot virtual disk and a target virtual disk. The virtual disk for which the point-in-time image is created (the source virtual disk) must be a standard virtual or thin disk in the storage array.

A snapshot virtual disk and a snapshot repository virtual disk are created during the online copy operation. The snapshot virtual disk is not an actual virtual disk containing data; instead, it is a reference to the data contained on the virtual disk at a specific time. For each snapshot taken, a snapshot repository virtual disk is created to hold the copy-on-write data for the snapshot. The snapshot repository virtual disk is used only to manage the snapshot image.

Before a data block on the source virtual disk is modified, the contents of the block to be modified are copied to the snapshot repository virtual disk. Because the snapshot repository virtual disk stores copies of the original data in those data blocks, further changes to those data blocks write only to the source virtual disk.

NOTE: If the snapshot virtual disk that is used as the copy source is active, the source virtual disk performance degrades due to copy-on-write operations. When the copy is complete, the snapshot is disabled and the source virtual disk performance is restored. Although the snapshot is disabled, the repository infrastructure and copy relationship remain intact.

Creating A Virtual Disk Copy For An MSCS Shared Disk

To create a virtual disk copy for a Microsoft Cluster Server (MSCS) shared disk, create a snapshot of the virtual disk, and then use the snapshot virtual disk as the source for the virtual disk copy.

NOTE: An attempt to directly create a virtual disk copy for an MSCS shared disk, rather than using a snapshot virtual disk, fails with the following error: The operation cannot complete because the selected virtual disk is not a source virtual disk candidate.

NOTE: When creating a snapshot virtual disk, map the snapshot virtual disk to only one node in the cluster. Mapping the snapshot virtual disk to the host group or both nodes in the cluster may cause data corruption by allowing both nodes to concurrently access data.

Virtual Disk Read/Write Permissions

After the virtual disk copy is complete, the target virtual disk automatically becomes read-only to the hosts. The target virtual disk rejects read and write requests while the virtual disk copy operation has a status of Pending or In Progress or if the operation fails before completing the copy. Keep the target virtual disk Read-only enabled if you want to preserve the data on the target virtual disk for reasons such as the following:

- If you are using the target virtual disk for backup purposes.
- If you are using the data on the target virtual disk to copy back to the source virtual disk of a disabled or failed snapshot virtual disk.

If you decide not to preserve the data on the target virtual disk after the virtual disk copy is complete, change the write protection setting for the target virtual disk to Read/Write.

Virtual Disk Copy Restrictions

Before you perform any virtual disk copy tasks, understand and adhere to the restrictions listed in this section. The restrictions apply to the source virtual disk, the target virtual disk, and the storage array.

- While a virtual disk copy has a status of In Progress, Pending, or Failed, the source virtual disk is available for read I/O activity only. After the virtual disk copy is complete, read and write I/O activity to the source virtual disk are permitted.
- A virtual disk can be selected as a target virtual disk for only one virtual disk copy at a time.
- A virtual disk copy for any virtual disk cannot be mounted on the same host as the source virtual disk.
- Windows does not allow a drive letter to be assigned to a virtual disk copy.
- A virtual disk with a Failed status cannot be used as a source virtual disk or target virtual disk.
- A virtual disk with a Degraded status cannot be used as a target virtual disk.
- Snapshots created using older (legacy) premium feature versions cannot be managed using newer snapshot premium feature options. A virtual disk in a snapshot group cannot be a target for a virtual disk copy.

  NOTE: If you want to choose the base virtual disk of an older (legacy) snapshot virtual disk as your target virtual disk, you must first disable all snapshot (legacy) virtual disks that are associated with the base virtual disk.
- A virtual disk participating in a modification operation cannot be selected as a source virtual disk or target virtual disk. Modification operations include the following:
  - Capacity expansion
  - RAID-level migration
  - Segment sizing
  - Virtual disk expansion
  - Defragmenting a virtual disk

    NOTE: The following host preparation sections also apply when using the virtual disk copy feature through the CLI interface.

Creating A Virtual Disk Copy

CAUTION: Possible loss of data—Source virtual disks that are participating in a virtual disk copy are available for read I/O activity only while a virtual disk copy has a status of In Progress or Pending. Write requests are allowed after the virtual disk copy has completed. If the source virtual disk has been formatted with a journaling file system, any attempt to issue a read request to the source virtual disk may be rejected by the storage array, and an error message may appear. The journaling file system driver issues a write request before it attempts to issue the read request. The storage array rejects the write request, and the read request may not be issued due to the rejected write request. This condition may result in an error message appearing, which indicates that the source virtual disk is write protected. To prevent this issue from occurring, do not attempt to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy has a status of In Progress. Also, make sure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy has completed to prevent error messages from appearing.

The Virtual Disk Copy premium feature includes these items:

- The **Create Copy Wizard**, which assists in creating a virtual disk copy
- The **Copy Manager**, which monitors virtual disk copies after they are created

Setting Read/Write Permissions On Target Virtual Disk

To set read/write permissions on the target virtual disk:

1. In the AMW, click **Storage & Copy Services**.
2. Select **Copy Services → Virtual Disk Copy → Manage Copies.**
   The **Copy Manager** window is displayed.
3. Select one or more copy pairs in the table.
4. Perform one of these actions:
   - To enable Read-only permission, select **Change → Target Virtual Disk Permissions → Enable Read-Only.**

     NOTE: Write requests to the target virtual disk are rejected when the Read-only permission is enabled on the target virtual disk.
   - To disable Read-only permission, select **Change → Target Virtual Disk Permissions → Disable Read-Only.**

Before You Begin

A virtual disk copy fails all snapshot virtual disks that are associated with the target virtual disk, if any exist. If you select a source virtual disk of a snapshot virtual disk, you must disable all of the snapshot virtual disks that are associated with the source virtual disk before you can select it as a target virtual disk. Otherwise, the source virtual disk cannot be used as a target virtual disk.

A virtual disk copy overwrites data on the target virtual disk and automatically makes the target virtual disk read-only to hosts.

If eight virtual disk copies with a status of In Progress exist, any subsequent virtual disk copy has a status of Pending, which stays until one of the eight virtual disk copies completes.

Virtual Disk Copy And Modification Operations

If a modification operation is running on a source virtual disk or a target virtual disk, and the virtual disk copy has a status of In Progress, Pending, or Failed, the virtual disk copy does not take place. If a modification operation is running

on a source virtual disk or a target virtual disk after a virtual disk copy has been created, the modification operation must complete before the virtual disk copy can start. If a virtual disk copy has a status of In Progress, any modification operation does not take place.

Create Copy Wizard

The **Create Copy Wizard** guides you through:

- Selecting a source virtual disk from a list of available virtual disks
- Selecting a target virtual disk from a list of available virtual disks
- Setting the copy priority for the virtual disk copy

When you have completed the wizard dialogs, the virtual disk copy starts, and data is read from the source virtual disk and written to the target virtual disk.

Operation in Progress icons are displayed on the source virtual disk and the target virtual disk while the virtual disk copy has a status of In Progress or Pending.

Failed Virtual Disk Copy

A virtual disk copy can fail due to these conditions:

- A read error from the source virtual disk
- A write error to the target virtual disk
- A failure in the storage array that affects the source virtual disk or the target virtual disk

When the virtual disk copy fails, a critical event is logged in the Event Log, and a Needs Attention icon is displayed in the AMW. While a virtual disk copy has this status, the host has read-only access to the source virtual disk. Read requests from and write requests to the target virtual disk do not take place until the failure is corrected by using the Recovery Guru.

Preferred RAID Controller Module Ownership

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. If both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped, ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed.

Failed RAID Controller Module

You must manually change RAID controller module ownership to the alternate RAID controller module to allow the virtual disk copy to complete under all of these conditions:

- A virtual disk copy has a status of In Progress
- The preferred RAID controller module of the source virtual disk fails
- The ownership transfer does not occur automatically in the failover

Copy Manager

After you create a virtual disk copy by using the **Create Copy Wizard**, you can monitor the virtual disk copy through the **Copy Manager**. From the **Copy Manager**, a virtual disk copy may be re-copied, stopped, or removed. You can also modify the attributes, such as the copy priority and the target virtual disk Read-Only attribute. You can view the status of a virtual disk copy in the **Copy Manager**. Also, if you want to determine which virtual disks are involved in a virtual disk copy, you can use the **Copy Manager** or the storage array profile.

Copying The Virtual Disk

You can create a virtual disk copy by using the Create Copy Wizard. A virtual disk copy automatically makes the target virtual disk read-only to hosts. You might want to keep this attribute enabled to preserve the data on the target virtual disk. To prevent write-protected error messages from appearing, do not try to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy has a status of In Progress. Also, make sure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy has completed to prevent error messages from appearing.

To prevent write-protected error messages from appearing, do not try to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy has a status of In Progress. Also, make sure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy has completed to prevent error messages from appearing.

CAUTION: Possible loss of data access—A virtual disk copy overwrites data on the target virtual disk.

CAUTION: If you decide not to preserve the data on the target virtual disk after the virtual disk copy has completed, disable the Read-Only attribute for the target virtual disk. See Virtual Disk Read/Write Permissions for more information on enabling and disabling the Read-Only attribute for the target virtual disk.

To copy the virtual disk:

1. Stop all I/O activity to the source virtual disk and the target virtual disk.
2. Unmount any file systems on the source virtual disk and the target virtual disk.
3. In the AMW, select the **Storage & Copy Services** tab.
4. Under **Virtual Disks** area, select the source virtual disk that you want to use for the online copy.
5. Right click on the selected source virtual disk and select **Create** → **Virtual Disk Copy** in the pop-up menu.
   The **Select Copy Type** wizard is displayed.
6. Select a copy type and click **Next**.

   NOTE: If you select **Offline**, the source virtual disk is not available for any I/O when the copy operation is in progress.

   The **Select Target Virtual Disk** window is displayed.
7. Select the appropriate target virtual disk and click **Next**.
   The **Confirmation** window is displayed.
8. In the **Copy Priority** area, select the relevant copy priority and type yes to confirm.
9. Click **Finish**.
   The **Preview** window displays the summary of your selections.

   NOTE: Operation in Progress icons appear on the source virtual disk and the target virtual disk while the virtual disk copy has a status of In Progress or Pending.

Storage Array Performance During Virtual Disk Copy

The following factors contribute to the overall performance of the storage array:

- I/O activity
- Virtual disk RAID level
- Virtual disk configuration — Number of drives in the virtual disk groups
- Virtual disk type — Snapshot virtual disks may take more time to copy than standard virtual disks
- Snapshots created using older RAID controller firmware versions (legacy snapshots) will take longer to complete

During a virtual disk copy, resources for the storage array are diverted from processing I/O activity to completing a virtual disk copy. This affects the overall performance of the storage array. When you create a new virtual disk copy, you define the copy priority to determine how much controller processing time is diverted from I/O activity to a virtual disk copy operation.

Setting Copy Priority

You can use the Copy Manager to select the rate at which a virtual disk copy completes for a selected copy pair. You can change the copy priority for a copy pair at any of these times:

- Before the virtual disk copy begins
- While the virtual disk copy has a status of In Progress
- When you re-create a virtual disk copy

To set copy priority:

1. In the AMW, select the **Storage & Copy Services** tab and select **Copy Services → Virtual Disk Copy → Manage Copies.**
   The **Copy Manager** window is displayed.
2. In the table, select one or more copy pairs.
3. Select **Change → Copy Priority.**
   The **Change Copy Priority** window is displayed.
4. In the **Copy Priority** area, select the appropriate copy priority, depending on your system performance needs.

   NOTE: There are five copy priority rates available:

   - lowest
   - low
   - medium
   - high
   - highest

   If the copy priority is set at the lowest rate, I/O activity is prioritized, and the virtual disk copy takes longer.

Stopping A Virtual Disk Copy

You can stop a virtual disk copy operation that has an In Progress status, a Pending status, or a Failed status. Stopping a virtual disk copy that has a Failed status clears the Needs Attention status displayed for the storage array.

Keep these guidelines in mind when you stop a virtual disk copy:

- To use this option, select only one copy pair in the Copy Manager.
- When the virtual disk copy is stopped, all of the mapped hosts have write access to the source virtual disk. If data is written to the source virtual disk, the data on the target virtual disk no longer matches the data on the source virtual disk.

To stop a virtual disk copy, complete the following steps:

1. In the AMW, select the **Storage & Copy Services** tab and select **Copy Services → Virtual Disks → Manage Copies.**
   The **Copy Manager** window is displayed.
2. Select the copy pair in the table.
3. Select **Copy → Stop.**
4. Click **Yes**.

Recopying A Virtual Disk

You can recopy a virtual disk when you have stopped a virtual disk copy and you want to start it again or when a virtual disk copy has failed. The Recopy option overwrites existing data on the target virtual disk and makes the target virtual disk read-only to hosts. This option fails all snapshot virtual disks associated with the target virtual disk, if any exist.

Preparing Host Servers To Recopy A Virtual Disk

NOTE: Before you create a new copy of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk (and, if applicable, the target disk) to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer, to make sure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the virtual disk copy.

Before creating a new virtual disk copy for an existing copy pair, both the host server and the associated virtual disk you are recopying have to be in the proper state. Perform the following steps to prepare your host server and virtual disk:

1. Stop all I/O activity to the source and target virtual disk.
2. Using your Windows system, flush the cache to both the source and the target virtual disk (if mounted). At the host prompt, type: `SMrepassist -f <filename-identifier>` and press <Enter>.
   For more information, see SMrepassist Utility.
3. Click the **Summary** tab, then click **Storage & Copy Services** to ensure that the virtual disk is in Optimal or Disabled status.
4. Remove the drive letter(s) of the source and (if mounted) virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the virtual disk. If this is not done, the copy operation reports that it has completed successfully, but the copied data is not updated properly.
5. Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable virtual disk copies.

   NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Recopying The Virtual Disk

You can use the Copy Manager to create a new virtual disk copy for a selected source virtual disk and a target virtual disk. Use this option when you have stopped a virtual disk copy and want to start it again or when a virtual disk copy has failed or completed. The virtual disk copy starts over from the beginning.

Keep these guidelines in mind when re-copying a virtual disk:

- If hosts are mapped to the source virtual disk, the data that is copied to the target virtual disk when you perform the re-copy operation might have changed since the previous virtual disk copy was created.
- Select only one virtual disk copy in the **Copy Manager** dialog.

CAUTION: Possible loss of data—The re-copying operation overwrites existing data on the target virtual disk.

CAUTION: Possible loss of data access—While a virtual disk copy has a status of In Progress or Pending, source virtual disks are available for read I/O activity only. Write requests are allowed after the virtual disk copy has completed.

To recopy the virtual disk:

1. Stop all I/O to the source virtual disk and the target virtual disk.
2. Unmount any file systems on the source virtual disk and the target virtual disk.
3. In the AMW, select **Copy Services → Virtual Disk Copy → Manage Copies.**
   The **Copy Manager** window is displayed.
4. Select the copy pair in the table.
5. Select **Copy → Re-Copy.**
   The Re-Copy window is displayed.
6. Set the copy priority.
   There are five copy priority rates available: lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized, and the virtual disk copy takes longer. If the copy priority is set to the highest priority rate, the virtual disk copy is prioritized, but I/O activity for the storage array might be affected.

Removing Copy Pairs

You can remove one or more virtual disk copies by using the **Remove Copy Pairs** option. Any virtual disk copy-related information for the source virtual disk and the target virtual disk is removed from the **Virtual Disk Properties** dialog and the **Storage Array Profile** dialogs. When you remove a virtual disk copy from the storage array, the Read-Only attribute for the target virtual disk is also removed. After the virtual disk copy is removed from the Copy Manager, you can either select the target virtual disk as a source virtual disk or the target virtual disk for a new virtual disk copy. If you remove a virtual disk copy, the source virtual disk and the target virtual disk no longer appear in the Copy Manager.

Keep these guidelines in mind when you remove copy pairs:

- Removing copy pairs does not delete the data on the source virtual disk or target virtual disk.
- If the virtual disk copy has a status of In Progress, you must stop the virtual disk copy before you can remove the copy pair.

To remove copy pairs:

1. In the AMW, select **Copy Services → Virtual Disk Copy → Manage Copies.**
   The **Copy Manager** window is displayed.
2. In the table, select one or more copy pairs.

3. Select **Copy → Remove Copy Pairs.**
   The **Remove Copy Pairs** dialog is displayed.
4. Click **Yes**.

13

Device Mapper Multipath For Linux

Overview

The MD Series storage arrays use a Linux operating system software framework, known as Device Mapper (DM), to enable multipath capabilities on Linux Host Servers. The DM multipath functionality is provided by a combination of drivers and utilities. This chapter describes how to use those utilities to complete the process of enabling MD Series storage arrays on a Linux system.

NOTE: The required Device Mapper software components are installed on a Linux host server by running the MD Series storage arrays resource DVD installation program on the server, and selecting either the Full or Host install option. For detailed installation procedures, see the storage array's Deployment Guide at **dell.com/support/manuals**.

Benefits of using DM Multipath include:

- Detects path failure and re-routes I/O to other available paths
- Revalidates failed paths after path restoration
- Utilizes multiple available paths to maximize performance
- Reconfigures path usage based on path states and error conditions
- Unifies multiple device nodes into a single logical multipath device node
- Indentifies a new multipathed LU and automatically configures a new multipath node
- Provides device name persistency for DM devices under **/dev/mapper/**

Using DM Multipathing Devices

NOTE: Using or modifying any nodes other than the multipathing device nodes can result in array or file system problems, including loss of communication with the array and corruption of the file system. Avoid accessing any device other than the multipathing device.

NOTE: After creating a partition on a multipathing device, all I/O operations, including file system creation, raw I/O and file system I/O, must be done through the partition node and not through the multipathing device nodes.

Prerequisites

The following tasks must be completed before proceeding. For more information about step 1 through step 3, see the storage array's Deployment Guide. For more information about step 4, see Creating Virtual Disks.

1. Install the host software from the MD Series storage arrays resource DVD — Insert the Resource media in the system to start the installation of Modular Disk Storage Manager (MD Storage Manager) and Modular Disk Configuration Utility (MDCU).

   NOTE: Installation of Red Hat 5.x requires a remount of the DVD media to make contents executable.

2. Reboot when prompted by the install program — The installation program prompts for and needs a reboot at completion of the installation.

3. Configure using MDCU — After the host server has rebooted, the MDCU automatically starts and is present on the desktop. This utility allows for quick and easy configuration of new and or existing MD Series storage arrays present on your network. It also provides a GUI Wizard for establishing the iSCSI sessions to the array.
4. Create and map virtual disks using the MD Storage Manager — After configuring the arrays using the MDCU, run the MD Storage Manager to create and map virtual disks.

Using The MD Storage Manager

Use the MD Storage Manager to:

- Map the host server to the MD Series storage array
- Create the virtual disks
- Map newly created arrays to your host server

NOTE: Any arrays configured with MDCU automatically get added to the list of devices in the EMW.

Device Mapper Configuration Steps

To complete the DM multipathing configuration and make storage available to the Linux host server:

1. Scan for virtual disks.
   See Scan For Newly Added Virtual Disks.
2. Display the multipath device topology.
   See Display The Multipath Device Topology Using The Multipath Command.
3. Create a partition on a multipath device node.
   See Create A New fdisk Partition On A Multipath Device Node.
4. Add a partition to DM.
   See Add A New Partition To Device Mapper.
5. Create a file system on a DM partition.
   See Create A File System On A Device Mapper Partition.
6. Mount a DM partition.
   See Mount A Device Mapper Partition.

The following instructions show how to complete each of these steps.

In the following command descriptions, <x> is used to indicate where a substitution must be made. On Red Hat Enterprise Linux systems, <x> is the number assigned to the device. On SUSE Linux Enterprise Server systems, <x> is the letter(s) assigned to the device.

Scan For Newly Added Virtual Disks

The `rescan_dm_devs` command scans the host server system looking for existing and newly added virtual disks mapped to the host server.

```
# rescan_dm_devs
```

If an array virtual disk (VD) is mapped to the host server at a later time the `rescan_dm_devices` command must be run again to make the VD a visible LUN to the operating system.

Display The Multipath Device Topology Using The Multipath Command

The multipath command adds newly scanned and mapped virtual disks to the Device Mapper tables and creates entries for them in the **/dev/mapper** directory on the host server. These devices are the same as any other block devices in the host.

To list all the multipath devices, run the following command:

```
# multipath –ll
```

The output must be similar to this example, which shows the output for one mapped virtual disk.

```
mpath1 (3600a0b80005ab177000017544a8d6b92) dm-0 DELL, MD3xxxx[size=5.0G]
[features=3 queue_if_no_path pg_init_retries 50][hwhandler=1 rdac][rw]\_ round-
robin 0 [prio=6][active] \_ 5:0:0:0   sdc  8:32   [active][ready]\_ round-robin
0 [prio=1][enabled] \_ 4:0:0:0   sdb  8:16   [active][ghost]
```

where:

`mpath1` is the name of the virtual device created by device mapper. It is located in the **/dev/mapper** directory.

`DELL` is the vendor of the device.

`MD3xxxx` is the model of the device.

`Sdc` is the physical path to the owning controller for the device.

`Sdb` is the physical path to the non-owning controller for the device.

The following is an example of SLES output:

```
mpathb(360080e500017b2f80000c6ca4a1d4ab8) dm-21 DELL,MD3xxxx[size=1.0G]
[features=3 queue_if_no_path pg_init_retries 50][hwhandler=1 rdac][rw]\_ round-
robin 0 [prio=6][active] \_ 4:0:0:22 sdx  65:112 [active][ready]\_ round-robin
0 [prio=1][enabled] \_ 6:0:0:22 sdcl 69:144 [active][ghost]
```

where:

`mpathb` is the name of the virtual device created by device mapper. It is located in the **/dev/mapper** directory.

`DELL` is the vendor of the device.

`MD3xxxx` is the model of the device.

`Sdx` is the physical path to the owning controller for the device.

`Sdcl` is the physical path to the non-owning controller for the device.

Create A New fdisk Partition On A Multipath Device Node

The `fdisk` command allows creation of partition space for a file system on the newly scanned and mapped virtual disks that have been presented to Device Mapper.

To create a partition with the multipathing device nodes **/dev/mapper/mpath<x>**, for example, use the following command:

```
# fdisk /dev/mapper/mpath<x>
```

where mpath<x> is the multipathing device node on which you want to create the partition.

NOTE: The <x> value is an alphanumeric operating system dependent format. The corresponding value for mapped virtual disks can be seen using the previously run multipath command. See your operating system documentation for additional information on fdisk.

Add A New Partition To Device Mapper

The `kpartx` command adds the new fdisk partition to the Device Mapper list of usable partitions. See examples below, where `mpath<x>` is the device node on which the partition was created.

```
# kpartx –a /dev/mapper/mpath<x>
```

If successful, the command does not display an output. To verify success and view exact partition naming, you can use these commands to see the full partition names assigned.

```
# cd /dev/mapper# ls
```

The following are some examples of the general mapping formats:

- On Red Hat Enterprise Linux (RHEL) hosts, a partition node has the format:**/dev/mapper/mpath<x>p<y>**
  Where <x> is the alphabetic number for the multipathing device, <y> is the partition number for this device.
- On SUSE Linux Enterprise Server (SLES) 11.x hosts, a partition node has the format:**/dev/mapper/mpath<x>-part<y>**
  Where <x> is letter(s) assigned to the multipathing device and <y> is the partition number.
- On SLES 10.3 hosts, a partition node has the format: **/dev/mapper/mpath<x>_part<y>**
  Where <x> is the letter(s) assigned to the multipathing device and <y> is the partition number.

NOTE: After creating a partition on a device capable of multipathing, all I/O operations, including file system creation, raw I/O and file system I/O, must be done through the partition node, and not through the multipathing device nodes.

Create A File System On A Device Mapper Partition

Use the standard `mkfs` command to create the file system on the newly created Device Mapper partition.

For example:

```
# mkfs –t <filesystem type> /dev/mapper/<partition node>
```

where `<partition node>` is the partition on which the file system is created.

Mount A Device Mapper Partition

Use the standard mount command to mount the Device Mapper partition, as shown below:

```
# mount /dev/mapper/<partition_node> <mounting point>
```

Ready For Use

The newly created virtual disks created on the MD Series storage array are now setup and ready to be used. Future reboots automatically find multipathing devices along with their partitions.

NOTE: To ensure data integrity protection, reboot a Linux host server attached to an MD Series storage array using the procedure given below.

Linux Host Server Reboot Best Practices

Its is recommended that you follow the procedures given below when you reboot your Linux host server using Device Mapper multipathing with an MD Series storage array.

1. Unmount all Device Mapper multipath device nodes mounted on the server: `# umount <mounted_multipath_device_node>`
2. Stop the Device Mapper multipath service: `# /etc/init.d/multipathd stop`
3. Flush the Device Mapper multipath maps list to remove any old or modified mappings: `# multipath –F`

   NOTE: The boot operating system drive may have an entry with the Device Mapper multipathing table. This is not affected by the `multipath –F` command.
4. Log out of all iSCSI sessions from the host server to the storage array: `# iscsiadm –m node --logout`

Important Information About Special Partitions

When using Device Mapper with the MD Series storage arrays, all physical disks are assigned a disk device node. This includes a special device type used for in-band management of the storage arrays, known as the Access Disk or Universal Xport device.

CAUTION: Certain commands, such as `lsscsi`, display one or more instances of Universal Xport devices. These device nodes must never be accessed, mounted, or used in any way. Doing so can cause loss of communication to the storage array and possibly cause serious damage to the storage array, potential making data stored on the array inaccessible.

Only multipathing device nodes and partition nodes created using the directions provided above must be mounted or in any way accessed by the host system or its users.

Table 7. Useful Device Mapper Commands

Command	Description
`multipath –h`	Prints usage information.
`multipath –ll`	Displays the current multipath topology using all available information (sysfs, the device mapper, path checkers, and so on).
`multipath`	Re-aggregates multipathing device with simplified output.
`multipath –f <multipath_dev_node>`	Flushes out Device Mapper for the specified multipathing device. Used if the underlying physical devices are deleted/unmapped.
`multipath –F`	Flushes out all unused multipathing device maps.
`rescan_dm_devs`	Dell provided script. Forces a rescan of the host SCSI bus and aggregates multipathing devices as needed. Use this command when: • LUNs are dynamically mapped to the hosts. • New targets are added to the host. • Failback of the storage array is required. • For MD Series iSCSI storage arrays, iSCSI sessions have to be established for rescan to take effect.

Limitations And Known Issues

- In certain error conditions with the `no_path_retry` or the `queue_if_no_path` feature is set, applications may hang. To overcome these conditions, enter the following command for each affected multipath device:

  `dmsetup message [device] 0 "fail_if_no_path"`

  where `[device]` is the multipath device name (for example, mpath2; do not specify the path)

- I/O may hang when a Device Mapper device is deleted before the volume is unmounted.
- If the `scsi_dh_rdac` module is not included in `initrd`, slower device discovery may be seen and the syslog may become populated with buffer I/O error messages.
- I/O may hang if the host server or storage array is rebooted while I/O is active. All I/O to the storage array should be stopped before shutting down or rebooting the host server or storage array.
- With an MD Series torage array, after a failed path is restored, failback does not occur automatically because the driver cannot auto-detect devices without a forced rescan. Run the command `rescan_dm_devs` to force a rescan of the host server. This restores the failed paths enabling failback to occur.
- Failback can be slow when the host system is experiencing heavy I/O. The problem is exacerbated if the host server is also experiencing very high processor utilization.
- The Device Mapper Multipath service can be slow when the host system is experiencing heavy I/O. The problem is exacerbated if the host server is also experiencing very high processor utilization.
- If the root disk is not blacklisted in the **multipath.conf** file, a multipathing node may be created for the root disk. The command `multipath -ll` lists vendor/product ID, which can help identify this issue.
- If upgrading from a previous version of SLES, uninstall and then reinstall the latest `scsi_dh_rdac` module on the updated SLES installation. Then update the kernel and install the MD Storage Manager from the DVD.

Troubleshooting

Question	Answer
How can I check if multipathd is running?	Run the following command. `/etc/init.d/multipathd status`
Why does the multipath –ll command output not show any devices?	First verify if the devices are discovered or not. The command `#cat /proc/scsi/scsi` displays all the devices that are already discovered. Then verify the **multipath.conf** to ensure that it is been updated with proper settings. After this, run `multipath`. Then run `multipath -ll`, the new devices must show up.
Why is a newly-mapped LUN not assigned a multipathing device node?	Run `rescan_dm_devs` in any directory. This should bring up the devices.
I removed a LUN. But the multipathing mapping is still available.	The multipathing device is still available after you remove the LUNs. Run `multipath -f <device node for the deleted LUN>` to remove the multipathing mapping. For example, if a device related with **/dev/dm-1** is deleted, you must run `multipath -f /dev/dm-1` to remove **/dev/dm-1** from DM mapping table. If multipathing daemon is stopped/restarted, run `multipath -F` to flush out all stale mappings.
Failback does not happen as expected with the array.	Sometimes the low level driver cannot auto-detect devices coming back with the array. Run `rescan_dm_devs` to rescan host server SCSI bus and re-aggregate devices at multipathing layer.

14

Configuring Asymmetric Logical Unit Access

If your MD Series RAID storage array supports Asymmetric Logical Unit Access (ALUA), active-active throughput allows I/O to pass from a RAID controller module to a virtual disk that is not owned by the controller. Without ALUA, the host multipath driver is required to send data requests targeted to a specific virtual disk to the owning RAID controller module. If the controller module does not own the virtual disk, it rejects the request.

ALUA Performance Considerations

While ALUA enables an MD-series storage array with a dual-controller (duplex) configuration to service I/O requests through either RAID controller module, performance is decreased when the non-owning RAID controller module accesses a virtual disk. To maintain the best possible throughput, the host driver communicates with the controller firmware to send data requests to the owning RAID controller, if possible.

Automatic Transfer Of Ownership

The RAID controller firmware automatically transfers virtual disk ownership if more than 75 percent of data I/O over the previous five minutes was routed to the non-owning RAID controller. This indicates that either the storage array has lost redundant connections or that some of the data paths to the virtual disk or disk group are not usable. MD Storage Manager launches the Recovery Guru (Virtual Disk Not on Preferred Path) if the condition still exists after the default alert five-minute delay time has expired. For more information, see Recovery Guru.

Native ALUA Support On Microsoft Windows And Linux

The following operating systems supported by your MD Series storage arrays also support ALUA natively:

- all supported Microsoft Windows operating systems
- Red Hat Enterprise Linux 6.2
- SUSE Linux Enterprise Server 11.2 with Service Pack 2

NOTE: No configuration steps are required to enable ALUA on the operating systems listed above.

Enabling ALUA On VMware ESX/ESXi

VMware ESX version 4.1 (update 1 and update 2) and ESX/ESXi 5.0 do not have Storage Array Type Plug-in (SATP) claim rules automatically set to support ALUA on the MD Series storage arrays. To enable ALUA, you must manually add the claim rule.

Manually Adding SATP Rule In ESX 4.1 (Update 1 And Update 2)

To manually add the SATP rule in ESX 4.1 (update 1 and update 2):

1. Run the following commands to delete the old claim rule and set the Dell-required rule:
   - `# esxcli nmp satp deleterule -s VMW_SATP_LSI -V DELL -M` *array_PID*

– `# esxcli nmp satp addrule -s VMW_SATP_ALUA -V DELL -M` *array_PID* `-c tpgs_off`

Where, *array_PID* is your storage array model/product ID. To select the appropriate *array_PID* for your storage array, see the following table.

Table 8. Storage Arrays and Model/Product IDs (*array_PID*)

Storage Array	*array_PID*
MD3260	`MD32xx`
MD3260i	`MD32xxi`
MD3660i	`MD36xxi`
MD3660f	`MD36xxf`

2. Reboot your ESX-based host server.

Manually Adding SATP Rule In ESX/ESXi 5.0

To manually add the SATP rule in ESX/ESXi 5.0:

1. Run the following command: `# esxcli storage nmp satp rule add -s VMW_SATP_ALUA -V DELL -M` *array_PID* `-c tpgs_on`

   Where, *array_PID* is your storage array model/product ID. To select the appropriate *array_PID* for your storage array, see the table Storage Arrays and Model/Product IDs (*array_PID*) in Manually Adding SATP Rule In ESX 4.1 (Update 1 And Update 2).
2. Reboot your ESX-based host server.

Verifying ALUA On VMware ESX/ESXi

To verify that the SATP claim rule you set is added in VMware ESX/ESXi:

- Run the following command for ESX 4.1 (all versions):

  `# esxcli nmp satp listrules -s VMW_SATP_LSI`

  Verify that the claim rule for the `VID/PID = Dell/`*array_PID* shows `Claim Options` with the `tpgs_off` flag specified.
- Run the following command for ESX/ESXi 5.0:

  `# esxcli storage nmp satp rule list -s VMW_SATP_ALUA`

  Verify that the claim rule for `VMW_SATP_ALUA` with the `VID/PID = Dell/`*array_PID* shows the `tpgs_on` flag specified.

Verifying If Host Server Is Using ALUA For MD Storage Array

To confirm that the host server is using the ALUA plugin:

- For ESX 4.1, run the following command: `#esxcli nmp device list`
- For ESX/ESXi 5.0, run the following command: `#esxcli storage nmp device list`

The value for **Storage Array Type** must be `VMW_SATP_ALUA` on each MD Series storage array.

Setting Round-Robin Load Balancing Policy On ESX/ESXi-Based Storage Arrays

NOTE: Perform this procedure after you have enabled ALUA on VMware ESX/ESXi and verified if the host server is using ALUA for the MD storage array. For more information, see Enabling ALUA On VMware ESX/ESXi and Verifying If Host Server Is Using ALUA For MD Storage Array.

To set a round-robin load balancing policy on your ESX/ESXi-based host server:

1. Do one of the following:
   - For ESX 4.x, run the following command:
     ```
     # esxcli nmp satp setdefaultpsp -psp VMW_PSP_RR -satp VMW_SATP_ALUA/
     VMW_SATP_LSI
     ```
   - For ESX/ESXi 5.x, run the following command:
     ```
     # esxcli storage nmp satp set --default-psp VMW_PSP_RR --satp
     VMW_SATP_ALUA/VMW_SATP_LSI
     ```
2. Reboot your ESX-based host server.

15

Premium Feature—Remote Replication

The following types of Remote Replication are supported on the MD storage array:

- Remote Replication — Standard asynchronous replication using point-in-time images to batch the resynchronization between the local and remote site. This type of replication is supported on both Fibre Channel and iSCSI storage arrays (not between).
- Remote Replication (Legacy) — Synchronous (or full-write) replication that synchronizes local and remote site data in real-time. This type of replication is supported on Fibre Channel storage arrays only.

NOTE: This chapter describes the Standard asynchronous replication premium feature only. To understand the Remote Replication (synchronous) premium feature, see Premium Feature- Remote Replication (Legacy).

About Asynchronous Remote Replication

Standard Remote Replication (asynchronous) is a premium feature that provides RAID controller-based data replication between a local and remote storage array on a per-virtual disk basis. By identifying primary (local) and secondary (remote) virtual disk pairs, called replicated pairs, write operations to the primary virtual disk of the pair are tracked by the RAID controller firmware and captured in a point-in-time image and transferred to the secondary virtual disk in the pair.

Remote Replication groups allow you to manage synchronization of both virtual disks to create a consistent data set across local and remote storage arrays. Point-in-time images on the primary virtual disk and the secondary virtual disk can be resynchronized in a batch approach that increases replication throughput. When data synchronization completes, the system uses the point-in-time images on the secondary virtual disk to ensure that the data is maintained in a consistent state during subsequent synchronization operations to the secondary virtual disk.

NOTE: The standard Remote Replication premium feature is supported on both iSCSI and Fibre Channel storage arrays.

Remote Replicated Pairs And Replication Repositories

Replicated pairs, comprising of a primary and secondary virtual disk, contain identical data copies as a result of data synchronization. Replication repository virtual disks are used to manage replication data synchronization and are required for both the primary virtual disk and secondary virtual disk in a replicated pair.

A replication repository consists of the following types of data:

- Resynchronization and recovery point images for both primary and secondary virtual disk.
- Log information that tracks regions on the primary virtual disk that is written between synchronization intervals. These logs are only used on the primary virtual disk, but are also written to the secondary virtual disk in case of a role reversal.
- Statistics for each replicated pair.

The replication repository is normally created automatically when you create a replicated pair. However, you can also create the repository manually.

Types Of Remote Replication

The following are the types of Remote Replication premium features supported on your storage array:

- Remote Replication — Also known as standard or asynchronous, it is supported on both iSCSI- and Fibre Channel-based storage arrays (both local and remote storage arrays must use the same data protocol) and requires a dual RAID controller configuration.
- Remote Replication (Legacy) — Also known as synchronous or full-write, it is supported on Fibre Channel storage arrays only.

Differences Between Remote Replication Features

As compared to the (synchronous) Remote Replication (Legacy) feature, the standard (asynchronous write) Remote Replication premium feature uses a point-in-time snapshot image to capture the state of the source virtual disk and only writes data that has changed since the last point-in-time image.

With standard Remote Replication, the remote storage array is not fully synchronized with the local storage array. As a result, in the event of a sudden, total loss of the remote storage array, some transactions could be lost.

With synchronous Remote Replication (Legacy), every data write to a source virtual disk is replicated to a remote virtual disk. This produces an identical, real-time remote of production data.

Other differences include:

- Number of repository virtual disk required — Standard Remote Replication requires a repository virtual disk to be created for each replicated pair (remote virtual disk-to-local virtual disk). Alternately, Remote Replication (Legacy) only requires a single repository virtual disk.
- Data protocol supported — Standard Remote Replication is supported on both iSCSI and Fibre Channel storage arrays. Remote Replication (Legacy) is supported only on Fibre Channel storage arrays.

  NOTE: Both remote and local storage arrays must be of the same data protocol -- replication between Fibre Channel and iSCSI storage arrays is not supported.
- Distance limitations — Distance between local and remote storage arrays is unlimited using the Standard Remote Replication premium feature. Remote Replication (Legacy) has a limitation of approximately 10 km (6.2 miles) between local and remote storage arrays, based on general latency and application performance requirements.

Examples Of Typical Use

Standard (asynchronous) Remote Replication is more network-efficient and generally more suitable in environments that require fast, non-stop processing. Remote backup consolidation, long-distance disaster recovery and 24 x 7 data protection are also common uses.

Synchronous Remote Replication (Legacy) is designed to provide replication between a relatively small number of local systems that require business continuity — for example, data center-type operations, local disaster recovery and other top-tier applications.

Upgrading To Asynchronous Remote Replication From Remote Replication (Legacy)

When you upgrade a RAID controller firmware version that supports both legacy and non-legacy Remote Replication premium features, all legacy Remote Replication configurations in the RAID controller remain unaffected and continue to function normally.

Remote Replication Requirements And Restrictions

To use the standard Remote Replication premium feature, you must have:

- Two storage arrays with write access and both these storage arrays must have sufficient space to replicate data between them.
- Each storage must have a dual-controller Fibre Channel or iSCSI configuration (single-controller configurations are not supported).
- Fibre Channel Connection Requirements — You must attach dedicated remote replication ports to a Fibre Channel fabric environment. In addition, these ports must support the Name Service.
- You can use a fabric configuration that is dedicated solely to the remote replication ports on each RAID controller module. In this case, host systems can connect to the storage arrays using fabric.
- Fibre Channel Arbitrated Loop (FC-AL), or point-to-point configurations, are not supported for array-to-array communications.
- Maximum distance between the local site and remote site is 10 km (6.2 miles), using single-mode fibre Gigabit interface converters (GBICs) and optical long-wave GBICs.
- iSCSI connection considerations:
  - iSCSI does not require dedicated ports for replication data traffic
  - iSCSI array-to-array communication must use a host-connected port (not the Ethernet management port).
  - The first port that successfully establishes an iSCSI connection is used for all subsequent communication with that remote storage array. If that connection subsequently fails, a new session is attempted using any available ports.

Restrictions On Using Remote Replication

- RAID level, caching parameters, and segment size can differ between replicated virtual disks.
- The secondary virtual disk must be at least as large as the primary virtual disk.
- Only standard virtual disks can be included in a replication relationship. Snapshot (Legacy) virtual disks, snapshot virtual disks, and thin virtual disks cannot be used.
- A primary virtual disk can be a source virtual disk or a target virtual disk in a virtual disk copy. A secondary virtual disk cannot be a source virtual disk or a target virtual disk unless a role reversal is initiated after the copy has completed. If a role reversal is initiated during a **Copy in Progress** status, the copy fails and cannot be restarted.
- A virtual disk can be included in only one replication relationship.
- A virtual disk participating in a copy request cannot be a replicated secondary virtual disk.

Setting Up Remote Replication

Setting up Remote Replication between local and remote storage arrays using MD Storage Manager consists of the following:

- Activating the Remote Replication premium feature on both the local and remote storage arrays
- Creating a remote Replication group on the local storage array
- Adding a replicated pair of virtual disks to the Remote Replication group

Activating Remote Replication Premium Features

Activating Remote Replication automatically reserves specific ports on each RAID controller module for data replication. After the port is reserved, any non-replication related I/O request to that port is rejected. Only RAID controller modules configured for Remote Replication can communicate with the reserved ports.

The Remote Replication premium feature must be activated on both the local and storage arrays.

NOTE: Perform the activation steps below on the local storage array first and then repeat them on the remote storage array.

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Select **Copy Services** → **Remote Replication** → **Activate.**
3. If both Remote Replication and Remote Replication (Legacy) premium features are supported on your storage array, select **Remote Replication**.

   NOTE: If you want to activate Remote Replication (Legacy), see Premium Feature- Remote Replication (Legacy) for activation and configuration details.
4. If you had selected standard Remote Replication, click **Finish**.

   The Premium feature activation is complete.
5. If you had selected Remote Replication (Legacy), in the **Create Repositories** window, select where the replication repository virtual disks for the Remote Replication (Legacy) feature must reside. Select one of the following:

   – Free capacity on existing disk pool or disk group — If this option is selected, a corresponding disk pool or disk group must be selected.
   – Unconfigured capacity on a new disk pool or disk group — If this option is selected, choose either **Disk Pool** or **Disk Group**.
   – Click **Next**.

   The **Create Disk Pool** wizard or the **Create Disk Group** wizard is displayed.
6. Click **OK**.

   The **Remote Replication Activated** window is displayed. The system performs when the Remote Replication premium feature is activated:

   – Logs out all hosts currently using the highest numbered Fibre Channel host port on the RAID controller modules.
   – Reserves the highest numbered Fibre Channel host port on the RAID controller modules for replication data transmissions.
   – Rejects all host communication to this RAID controller module host port as long as the replication feature is active.
   – If the Remote Replication (Legacy) feature has been activated, the two replication repositories are created.

   NOTE: Repeat these steps to activate the remote replication premium features on the remote storage array.

Deactivating Remote Replication

Deactivating the Remote Replication premium feature removes RAID controller module port restrictions.

NOTE: Before deactivating the Remote Replication premium feature, delete all existing Remote Replication groups and replicated virtual disk pairs from the local and remote storage arrays.

To deactivate the Remote Replication feature:

1. From the AMW, select **Copy Services → Remote Replication → Deactivate.**
   A message prompts you to confirm if the Remote Replication premium feature is to be deactivated.
2. Click **Yes**.

Remote Replication Groups

After the Remote Replication premium feature is successfully activated on both the local and remote storage arrays, you can create a Remote Replication group on the local storage array.

This group will contain at least one replicated virtual disk pair—one on the local storage and one on the remote storage array. These disks serve as primary and secondary disks that share data synchronization settings to provide consistent backup between both storage arrays. Multiple replicated pairs can reside in a Remote Replication group, but each pair can only be a member of one Remote Replication group. For more information, see Remote Replication Group Requirements And Guidelines.

Purpose Of A Remote Replication Group

By creating a Remote Replication group, all replication virtual disk pairs in the group can be managed as one. For example, all replicated virtual disk pairs in a group can share the same data synchronization settings, primary and secondary roles, and write modes.

The following attributes also apply to a Remote Replication group:

- The local storage array serves as the primary side of the Remote Replication group, while the remote storage array serves as the secondary side of the Remote Replication group.
- At the virtual disk level, all virtual disks added to the Remote Replication group on the local storage array serve as the primary role in the Remote Replication configuration. Virtual disks added to the group on the remote storage array serve the secondary role.

Because applications may use more than one virtual disk, Remote Replication groups must be replicated as a pair. All members of the Remote Replication group are synchronized as a coordinated data set to provide consistent backup to the remote site.

Remote Replication Group Requirements And Guidelines

- The Remote Replication premium feature must be enabled and activated on the local and remote storage arrays used in the replication configuration.
- Both local and remote storage arrays must be connected through a supported Fibre Channel or iSCSI connection.
- The remote storage array must contain a virtual disk with a capacity greater than or equal to the capacity of the virtual disk you intend to include as its pair on the local storage array.
- By default, any new Remote Replication group is created empty:
  - Only replicated pairs can be added to a Remote Replication group.
  - Each replicated pair can be a member of only Remote Replication group.
- An unnamed storage array will be displayed in the Remote Replication Repository view in MD Storage Manager and labeled as unnamed.

Creating A Remote Replication Group

NOTE: The **Create Remote Replication Group** option is available on the local storage array only. A Remote Replication group cannot be created on the remote storage array.

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Select **Copy Services → Remote Replication → Remote Replication → Replication Group → Create.**
   The **Create Remote Replication Group** window is displayed.
3. In **Remote replication group name**, enter a group name (30 characters maximum).
4. In the **Choose the remote storage array** drop-down, select a remote storage array.

   NOTE: If a remote storage array is not available, you cannot continue. Verify your network configuration or contact your network administrator.
5. In the **Connection type** drop-down, choose your data protocol (iSCSI or Fibre Channel only).
6. Select **View synchronization settings** to set the synchronization settings for your Remote Replication group.
   For more information on the synchronization settings, see the online help topics.
7. Click **OK**.
   The Remote Replication group is created.

Replicated Pairs

The last step in setting up Remote Replication is creating a replicated pair of virtual disks and placing them in an already-created Remote Replication group.

A replicated pair consists of two virtual disks, one serving as the primary virtual disk on the local storage array and the other serving as the secondary virtual disk on the remote storage array. In a successful Remote Replication configuration, both these virtual disks contain identical copies of the same data. The replicated pair is contained in Remote Replication group, allowing them to synchronize at the same time as any other replicated pairs within the same Remote Replication group.

At the I/O level, all write operations are performed first to the primary virtual disk and then to the secondary virtual disk.

Guidelines for Choosing Virtual Disks in a Replicated Pair

The first step of creating a replicated pair begins by adding a virtual disk to the Remote Replication group on the local storage array. This virtual disk then becomes the primary virtual disk in the remote replicated pair. When a virtual disk on the remote storage array is added to same Remote Replication group, the replicated pair creation process is complete. This remote storage virtual disk becomes the secondary virtual disk in the replicated pair.

The two virtual disks -- one on the local storage array and one on the remote storage array -- essentially function as a single entity and allow you to manage the pair in tandem, not as two individual virtual disks.

Guidelines For Choosing Virtual Disks In A Replicated Pair

The following guidelines apply:

- Only standard virtual disks can be used in a replicated pair. Thin provisioned or snapshot virtual disks (any type) cannot be used.
- The Remote Replication premium feature must be enabled and activated on the local and remote storage arrays used for replication before creating replication pairs or Remote Replication groups.
- Local and remote storage arrays must be connected using supported Fibre Channel or iSCSI connections.
- The remote storage array must contain a virtual disk that is greater than or equal to the capacity of the primary virtual disk on the local storage array.
- Creating a replicated pair requires you to use the AMW of the local storage array and the AMW of the remote storage array to complete the creation process. Make sure that you have access to both storage arrays.

Creating Replicated Pairs

This procedure describes how to create the remote replicated pair on an existing remote replication group. To create a new Remote Replication group, see Creating a Remote Replication Group.

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Select **Copy Services → Remote Replication → Remote Replication → Replication Group → Create Replication Pair.**

   The **Select Remote Replication Group** window is displayed.

   NOTE: If the local storage array does not contain any Remote Replication groups, you must create one on the local storage array before proceeding.
3. Select an existing Remote Replication group, then click **Next**.
4. In the **Select Primary Virtual Disk** window, select one of the following:

   – Select an existing virtual disk on the local storage array to serve as the primary virtual disk in the replicated pair and click **Next**. Go to step 4.
   – Select the option to create a new virtual disk and click **Next**. See Creating a Standard Virtual Disk.
5. In the **Select Repository** window, select whether you want to create the replication repository automatically or manually:

   – Automatic — Select **Automatic** and click **Finish** to create the replication repository with default capacity settings.
   – Manual — Select **Manual** and click **Next** to define the properties for the replication repository. Then click **Finish**.

   NOTE: The replication repository is normally created automatically during virtual disk pair creation. Manual repository creation is recommended only for advanced storage administrators who understand physical disk redundancy and optimal physical disk configurations. The Automatic method is recommended.
6. Click **OK** when you see a message that the pair is successfully created.

Creating Replicated Pairs On The Remote Storage Array

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Select **Copy Services → Remote Replication → Remote Replication → Replication Group → Complete Replication Pair.**

   The **Complete Remote Replicated Pair** window is displayed.
3. Do one of the following:

   – Select **Automatic** and select an existing disk pool or disk group from the table, then click **Finish** to automatically complete the replicated pair creation process with the default secondary virtual disk selection and repository settings.
   – Select **Manual**, then click **Next** to choose an existing virtual disk as the secondary virtual disk and define the repository parameters for the remote side of the remote replicated pair.

   The Remote Replicated pair is created.

   The following occurs:

   – Initial synchronization between the local storage array and the remote storage array automatically begins.
   – The replicated pair and its properties are displayed under the individual virtual disk node for the secondary virtual disk.

– The **Associated Replicated Pairs** table is updated to show the replication information for the Remote Replication group.

Removing A Replicated Pair From A Remote Replication Group

Removing a replicated pair from a Remote Replication group breaks the replication relationship between the primary virtual disk on the local storage array and the secondary virtual disk on the remote storage array. Data on the virtual disks is not affected. As a result of this operation, the primary virtual disk and the secondary virtual disk become standard, host-accessible, non-replicated virtual disks.

When you remove a replicated pair from a Remote Replication group, the replication relationship is first removed on the local storage array, then from the remote storage array.

NOTE: Occasionally, when the removal process fails to complete on both storage arrays, the next data synchronization initiated by the primary virtual disk to the secondary virtual disk is paused. The Logical view in the AMW may also show an unresponsive secondary virtual disk. Removing the replication relationship on the local storage array must correct the problem.

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Select the Remote Replication group containing the replicated pair you want to remove and select one of the following:
   – **Copy Services → Remote Replication → Remote Replication → Replication Group → Remove.**
   – From the **Associated Replicated Pairs** table in the right pane, select the replicated pair you want to remove and select **Copy Services → Remote Replication → Remote Replication → Replication Pair → Remove.**

   The **Confirm Remove Replicated Pair** window is displayed.
3. Type `yes` and click **Remove**.

   NOTE: When you remove a replicated pair, the system deletes the associated replication repositories. To preserve them, de-select **Delete replicated pair repositories**.

16
Premium Feature—Remote Replication (Legacy)

NOTE: The following section applies to the Fibre Channel-only Remote Replication premium feature. This premium feature is referred to as Legacy to differentiate from the Remote Replication premium feature available on both Fibre Channel and iSCSI-based MD storage arrays. For information, see Premium Feature - Remote Replication.

The following types of Remote Replication are supported on the MD storage array:

- Remote Replication — Standard asynchronous replication using point-in-time images to batch the resynchronization between the local and remote site. This type of replication is supported on both Fibre Channel and iSCSI storage arrays (not between).
- Remote Replication (Legacy) — Synchronous (or full-write) replication that synchronizes local and remote site data in real-time. This type of replication is supported on Fibre Channel storage arrays only.

NOTE: This chapter describes the Remote Replication (Legacy) premium feature only. To understand the standard Remote Replication (asynchronous) premium feature, see Premium Feature - Remote Replication.

Overview

The Remote Replication (Legacy) premium feature enables real-time replication of data between two Fibre Channel storage arrays in separate locations. Using this type of remote replication, up to 16 replicated virtual disk pairs can be created on a storage array.

A replicated virtual disk pair is created from two standard virtual disks, both of which are logical data storage structures created on a storage array. A standard virtual disk can only be a member of one replicated pair. A pair consists of a primary virtual disk at a local storage array and a secondary virtual disk at a remote storage array. Data written by the host to the primary virtual disk is replicated directly from the primary storage array to the secondary virtual disk. No host server or application resources are used during replication. The data can be written either synchronously or asynchronously.

In the event of a disaster, loss of communications or catastrophic failure in the local storage array, the secondary virtual disk in the remote storage array can be promoted to the role of primary virtual disk and assume responsibility for maintaining business operations.

Switchable Host Access Configuration Required With Remote Replication (Legacy)

Access to the host server from the local and/or remote storage arrays using remote replication (Legacy) must be established through a supported Fibre Channel switch configuration. Direct-attached configurations between the host server and Dell PowerVault MD storage arrays are not supported.

For more information on configuring the switch environment, see Required Switch Zoning Configurations.

Activating Remote Replication (Legacy)

NOTE: If you ordered the Remote Replication (Legacy) feature, you received a Premium Feature Activation card in the Dell PowerVault MD Series storage array shipping box. Follow the directions on the card to install and enable this feature.

After the Remote Replication (Legacy) feature has been installed following the steps on the Premium Feature Activation card, it must be activated using a wizard based process. As part of activation, the following actions occur:

- Two replication repository virtual disks are created using existing capacity on either a disk group or storage array specified by the user.
- All hosts currently using the highest-numbered Fibre Channel host port (Port 3) on each RAID controller are logged out.
- The highest-numbered Fibre Channel host port (Port 3) on each RAID controller is reserved for replication data transmission.
- The highest-numbered Fibre Channel host port (Port 3) on each RAID controller is not available to direct host access when the Remote Replication premium feature is activated. Any attempted host communication to this port is rejected.
- Communication between the replication-reserved Fibre Channel ports requires established zoning of Fibre Channel switches in the fabric, see Required Switch Zoning Configurations.
- Before establishing virtual disk replication pairs, the Remote Replication premium (Legacy) feature must be activated on each of the storage arrays participating in replication.

Replication Repository Virtual Disks

Using the Remote Replication (Legacy) premium feature on the storage array requires that two replication repository virtual disks be created on one of the disk groups on the storage array. The RAID controller module stores replication information on these virtual disks, including information about incomplete remote writes that can be used to recover from RAID controller module resets and accidental storage array shutdown.

Other information about replication repository virtual disks include:

- Specifying that replication repository virtual disks be created from unconfigured free capacity on the disk group, or creating a new disk group and its member replication repository virtual disks from unconfigured free capacity on the storage array.

  NOTE: The disk group containing replication repository virtual disks cannot be deleted without disabling the Remote Replication premium feature.
- Default names of the replication repository virtual disks are **Replication Repository 1** and **Replication Repository 2**. These names cannot be changed.
- The activation process creates replication repository virtual disk with identical capacity. In a dual RAID controller module storage array, the default capacity for each replication repository virtual disks is 128 MB per controller (requiring a total of 256 MB in the RAID group). This default capacity cannot be modified.

RAID Levels For Replication Repository Virtual Disks

If replication repository virtual disks are created from unconfigured free capacity on the storage array, you can select a RAID level for the disk group of either RAID 1/10, RAID 5, or RAID 6. However, if the replication repository virtual disks are created from an existing disk group, RAID levels are not user configurable.

Primary And Secondary Virtual Disk Pairs

After the Remote Replication (Legacy) feature is activated and configuration of the FC fabric, individual virtual disk replication pairs are be created. Up to 16 replication virtual disk pairs can be established. Each pair consists of a primary

(online) virtual disk and a secondary (offline) virtual disk, both of identical capacity. During operation, primary and secondary roles can be interchanged to test the replicated data sets without breaking the relationship. Independent copies may be created by deleting the relationship, in which case both virtual disks are available online to their configured hosts.

Before creating a virtual disk pair, verify the following:

- The Remote Replication premium feature has been installed, enabled and activated on the primary and secondary storage arrays.
- The storage arrays containing the two virtual disks you want to replicate are connected through a Fibre Channel fabric interface.
- The secondary storage array contains a virtual disk that is equal (or larger) in capacity than the primary virtual disk in the replicated pair.
- Both virtual disks are visible to the MD Storage Manager.
- All I/O into the secondary virtual disk is stopped and all file systems are unmounted.

CAUTION: Creating a replicated virtual disk pair starts a synchronous process between the primary and secondary virtual disks. This process overwrites all existing data on the secondary virtual disk and sets the disk to read-only access. If you have existing data on your secondary virtual disk, back up the data before creating the replicated virtual disk pair.

Using Remote Replication (Legacy) With Other Features

Remote replication can be used with other disk utility features such as:

- Storage Partitioning
- Snapshot Virtual Disk Premium Feature
- Virtual Disk Copy Premium Feature
- Virtual Disk Expansion (VDE)

Snapshot Virtual Disk and Virtual Disk Copy are premium features that must be enabled and activated on the primary storage array.

Storage Partitioning With Remote Replication (Legacy)

Storage partitioning allows hosts to share access to virtual disks in a storage array. A storage partition is created when a collection of hosts (host group) or a single host is defined, then a virtual disk-to-logical unit number (LUN) mapping is defined. This mapping allows you to define which host or host group has access to a particular virtual disk in the storage array.

Storage partition definitions for the local storage array and remote storage array are independent. Establishing similar definitions on the disk serving in a secondary role reduces administrative effort required during site recovery.

Snapshot Virtual Disk With Remote Replication (Legacy)

A snapshot virtual disk is a point-in-time image of a virtual disk. In a remote replication, do not mount a snapshot virtual disk on the same server as the primary virtual disk.

Virtual Disk Copy With Remote Replication (Legacy)

The virtual disk copy premium feature copies data from a source virtual disk to a target virtual disk within the same storage array.

A primary virtual disk in a remote replication can be either a source virtual disk or a target virtual disk in a virtual disk copy.

You can create a virtual disk copy on the primary virtual disk in a replicated pair, but you cannot create a virtual disk copy of a secondary virtual disk in a replicated pair. Instead, you can copy a secondary virtual disk in one of two ways:

- Promote the secondary virtual disk to the primary role.

  CAUTION: If a primary/secondary role reversal is begun while a virtual disk copy is in progress, the virtual disk copy fails and be unable to restart.
- Create a snapshot virtual disk of the secondary virtual disk, then perform a virtual disk copy of the snapshot virtual disk.

Virtual Disk Expansion With Remote Replication (Legacy)

Virtual Disk Expansion (VDE) increases the capacity of an existing virtual disk. This increased capacity is achieved by using the free capacity available on one of the following:

- The disk group of the standard virtual disk
- The snapshot repository virtual disk

Performing a VDE operation does not interrupt access to data on disk groups, virtual disks or physical disks.

A VDE operation can be performed on a primary virtual disk or a secondary virtual disk of a replicated pair. It cannot be performed on a replication repository virtual disk. No other action can be performed on a virtual disk while it is undergoing expansion (except normal disk I/O).

NOTE: To perform a VDE operation, the remote replication must be in Optimal status. The properties pane in the logical view displays the status of a virtual disk.

Required Switch Zoning Configurations

Due to potential restrictions at the host level, remote replication configurations must contain Fibre Channel switches. Fibre Channel switches must be zoned so that a single host adapter accesses only one RAID controller module in a storage array. Additionally, all zone configurations should specify separate zones for ports that are reserved for remote replication functions.

NOTE: Do not zone the uplink port (E_port) that connects (cascades) switches within a fabric.

Switch zoning configurations are typically set using the switch management software provided by the Fibre Channel switch vendor. This software is either included with the switch itself, or downloaded from the vendor's website.

Zoning Guidelines for Remote Replication (Legacy)

Enabling remote replication on your storage array automatically dedicates one port per RAID controller exclusively to replication. These dedicated ports, Port 3 on RAID controller 0 and Port 3 on RAID controller 1, must be in different zones on the Fibre Channel switch. Generally, switch zoning should be established so that:

- All RAID Controller 0 Fibre Channel Port 3s are in a common zone.
- All RAID Controller 1 Fibre Channel Port 3s are in a common zone, but separately zoned from RAID Controller 0 ports.
- Host Fibre Channel access ports are in separate zones from those established for remote replication.

Switch Cascading

When two or more Fibre Channel switches are cascaded, the switch management software combines the ports for all of the switches that are linked.

Journaling File Systems And Remote Replication (Legacy)

When using a journaling file system, you cannot gain read-only access to a remote virtual disk. A journaling file system does not let you mount the remote virtual disk in Windows NTFS. However, you can mount a snapshot of the remote virtual disk, if available.

Prerequisites For Setting Up And Managing Remote Replication (Legacy)

Before setting up remote replication between two storage arrays, ensure that:

- The Remote Replication premium feature has been activated.
- The local storage array contains two replication repository virtual disks (created during activation).
- The local storage array contains the primary virtual disk to be replicated.
- The remote storage array contains the secondary virtual disk to serve as the replication target, and:
  - Both virtual disks are visible to a single management station; replicated disk pairs cannot be managed on separate instances of the MD Storage Manager.
  - The RAID level of the secondary virtual disk can be different than that of the primary virtual disk.
  - The capacity of the secondary virtual disk (after creation and formatting) must be equal to or greater than the capacity of the primary virtual disk.
- The storage arrays containing the virtual disks you want to replicate are connected to each other through a Fibre Channel switch fabric and are accessible from a single management station.

If either the primary or secondary virtual disk does not exist, you must create it before proceeding with remote replication setup. See Activating The Remote Replication Premium Feature And Creating Replication Virtual Disks and complete the necessary steps.

If all virtual disks are created, skip to Creating A Remote Replication.

Activating The Remote Replication (Legacy) Premium Feature And Creating Replication Virtual Disks

To activate the Remote Replication (Legacy) feature and create the two replication repository virtual disks required, you can either:

- Activate the feature, then create the two replication virtual disks from the total unconfigured capacity on the storage array, or
- Activate the feature, then create the two replication virtual disks from the unconfigured capacity in an already-existing disk group

Activating And Creating Replication Repository Virtual Disks From The Storage Array

To activate the Remote Replication feature and create the two replication repository virtual disks:

1. From the AMW, select **Copy Services → Remote Replication → Activate.**
2. In the Activate Remote Replication wizard, select **Unconfigured capacity (create a new disk group)** and click **Next**.

   The **Activate Remote Replication - Create Disk Group** wizard is displayed.
3. In **Disk Group Name**, type a unique name for the disk group.
4. Select one of the following to select a physical disk:

   – **Automatic** — The Storage Manager generates a list of available capacity and physical disk options for each available RAID level.

   – **Manual** — The Storage Manager generates a list of unselected physical disks.
5. Click **Next**.

   – If you selected **Automatic**, an empty **Select Capacity** table and a drop-down list of available RAID levels is displayed. Go to Step 5.

   – If you selected **Manual**, a populated **Unselected Physical Disks** table, empty **Selected Physical Disks** table, and a drop-down list of available RAID levels is displayed.
6. In the **Select RAID level** drop-down list, select the RAID level for the disk group. The **Select capacity** table displays the available virtual disks for the RAID level.
7. In the **Select capacity** table, select the physical disks and capacities for the new disk group, then click **Next**.

   The **Preview (Activate Remote Replication)** wizard is displayed.
8. Click **Finish**.

   The `Completed (Activate Remote Replication)` message is displayed.
9. Click **OK**.

   The Remote Replication premium feature is now active. The object tree displays the new disk group and the two replication repository virtual disks.

Activating And Creating Replication Repository Virtual Disks From An Existing Disk Group

Using this method for creating the replication repository virtual disks, free capacity in an existing disk group is used. By default, the replication repository virtual disks each have either 128 MB or 256 MB capacity. You cannot create the replication repository virtual disks on a disk group with insufficient capacity, or change the default capacities of the replication repository virtual disks.

1. In the menu bar in the AMW, select **Storage → Virtual Disk → Remote Replication → Create.**

   The **Introduction (Activate Remote Replication)** wizard is displayed.
2. Select **Free capacity on existing disk groups**.
3. From the list of available disk groups, select a disk group to contain the replication repository virtual disks, then click **Next**.

   The **Preview (Activate Remote Replication)** wizard is displayed.
4. Click **Finish**.

   The `Completed (Activate Remote Replication)` message is displayed.
5. Click **OK**.

   The Remote Replication premium feature is now active. The Logical pane displays the two replication repository virtual disks as part of the disk group.

Creating A Remote Replication (Legacy)

Before creating a remote replication, ensure that all prerequisites are met.

1. Open the AMW of both the local and remote storage array.
2. Verify that the Remote Replication premium feature has been activated on both storage arrays.
3. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
4. In the Logical pane of the local storage array, select the virtual disk you created to serve as your primary virtual disk. This is the disk that is replicated to the secondary disk.
5. In the menu bar in the AMW, select **Storage → Virtual Disk → Remote Replication → Create.**

   The **Introduction (Create Remote Replication)** wizard is displayed.
6. Click **Next**.

   The **Select Storage Array (Create Remote Replication (Legacy)** dialog is displayed. The **Storage Arrays** list shows the remote storage arrays.
7. Select the storage array where you created the secondary virtual disk, then click **Next**.

   The **Select Secondary Virtual Disk (Create Remote Replication (Legacy)** wizard is displayed. Go to Selecting The Secondary Virtual Disk.

Selecting The Secondary Virtual Disk

CAUTION: Creating a replicated virtual disk pair starts a process between the primary and secondary virtual disks that overwrites all existing data on the secondary virtual disk and set the disk to read-only access. If you have existing data on your secondary virtual disk stop all I/O to the disk, back up the data, and unmount any file systems mounted to the secondary virtual disk before creating the replicated virtual disk pair.

1. Select the remote storage array and the primary virtual disk.
2. In the **Select Secondary Virtual Disk (Create Remote Replication (Legacy)** wizard, select the secondary virtual disk.

   The secondary virtual disk must have a capacity equal to or greater than the capacity of the primary virtual disk.
3. Click **Next**.

   The **Set Write Mode (Create Remote Replication (Legacy)** wizard is displayed. Go to Setting The Write Mode.

Setting The Write Mode

The secondary host ports on the storage arrays are reserved for data synchronization between the primary virtual disk and the secondary virtual disk in a replicated virtual disk pair. You can set remote replication to write either synchronously or asynchronously.

- In the synchronous mode, the RAID controller module on the primary virtual disk sends an I/O completion message back to the host storage array after the data has been successfully copied to the secondary storage array. The synchronous mode is the preferred mode of operation, since it offers the best chance of full data recovery from the secondary storage array in the event of a disaster. However, a synchronous replication mode can degrade the I/O performance of the host when long-distance data transfer is involved.
- In the asynchronous mode, the RAID controller module on the primary storage array sends an I/O completion message to the host storage array before the data is successfully copied to the secondary storage array. The asynchronous mode offers faster host I/O performance; however, it does not guarantee that data was successfully written to the secondary virtual disk or that the write requests were completed on the secondary virtual disk in the same order they were initiated. In asynchronous mode, you can also specify whether to use a write consistency group option. A write consistency group ensures that the secondary virtual disk receives write requests in the sequence initiated by the RAID controller module of the primary virtual disk. The secondary virtual disk can also be added to a write consistency group.

To set the write mode for the remote replication:

1. In the **Set Write Mode (Create Remote Replication (Legacy)** wizard, select either the **Synchronous** or **Asynchronous** mode.
2. If you select asynchronous mode, select whether to use the **Add to write consistency group** option.

3. Click **Next**.

   The **Select Synchronization Settings (Create Remote Replication (Legacy)** wizard is displayed. Go to Setting Synchronization Priority And Synchronization Method.

Setting Synchronization Priority And Synchronization Method

You can set the priority for allocating system resources to synchronizing the remote replication.

- Higher synchronization priorities allocate more resources to the process and might degrade I/O performance.
- Lower synchronization priorities allocate fewer resources to the process and have less impact on normal I/O performance.

The initial synchronization priority and synchronization method you select can be changed later. For more information about resynchronizing virtual disks in a remote replication, go to Resynchronizing Virtual Disks.

To set Synchronization Priority and Synchronization Method:

1. In the **Select Synchronization Settings (Create Remote Replication (Legacy)** wizard, select the synchronization priority on the slide bar.
2. Select either **Manual** or **Automatic** resynchronization.
   - Automatic resynchronization — Resynchronization starts immediately after communication is restored between unsynchronized replicated virtual disks.
   - Manual resynchronization (recommended) — The replicated pair must be manually resynchronized each time communication is restored between unsynchronized replicated virtual disks.
3. Click **Next**.

   The **Preview (Create Remote Replication(Legacy)** wizard is displayed. Go to Completing The Remote Replication.

Completing The Remote Replication (Legacy)

CAUTION: Creating a replicated virtual disk pair starts a process between the primary and secondary virtual disks that overwrites ALL existing data on the secondary virtual disk and set the disk to read-only access. If you have existing data on your secondary virtual disk: (1) stop all I/O to the disk, (2) back up the data and (3) unmount any file systems mounted to the secondary virtual disk before creating the replicated virtual disk pair.

After selecting the synchronization settings, perform these steps to complete the remote replication:

1. In the **Preview (Create Remote Replication (Legacy)** wizard, type `Yes` to confirm the synchronization and click **Finish**.

   The following confirmation message is displayed:

   ```
   Creation Successful (Create Remote Replication)
   ```
2. If you want to create another replication pair, click **Yes** and repeat the steps. Otherwise, click **No** to exit.
   - In the AMWs of both the local storage array and the remote storage array, the Logical panes show the replicated virtual disk pairs as members of their disk groups.
   - In the local storage array, the Properties pane displays the Replication status as Synchronizing and the Synchronization – Progress bar shows estimated time to completion.

   To view detailed information about the virtual disks in a remote replication, go to either Viewing Information About A Remote Replication Or Replication Repository Virtual Disk. Each virtual disk is also represented by a status icon in the **Devices** tab on the EMW. For a description of each icon, see Virtual Disk Status Icons.

RAID Controller Module Ownership/Preferred Path

During a remote replication operation, the RAID controller module that owns the primary virtual disk must correspond to the RAID controller owning the secondary volume on the remote array. If both virtual disks do not have the same preferred RAID controller module when a remote replication begins, the ownership of the secondary virtual disk is automatically transferred to the preferred RAID controller module of the primary virtual disk.

- When the remote replication is completed or is stopped, ownership of the secondary virtual disk is restored to its preferred RAID controller module.
- If ownership of the primary virtual disk is changed during the remote replication, ownership of the secondary virtual disk is also changed.

If any of the following conditions apply to your RAID controller module, you must manually change module ownership to the alternate RAID controller module to allow the remote replication to finish:

- The status of the remote replication is **In Progress**.
- The preferred RAID controller module of the primary virtual disk fails.
- Transfer of ownership does not occur automatically during failover.

CAUTION: To avoid possible loss of data, verify that either the virtual disks are not in use, or a multi-path driver is installed on the host. If you change the RAID controller module ownership/preferred path while an application is using one of the virtual disks, I/O activity is disrupted and I/O errors may occur.

Viewing Information About A Remote Replication (Legacy) Or Replication Repository (Legacy) Virtual Disk

In MD Storage Manager, the Storage Array Profile and the Properties pane display information about the physical and logical characteristics of the remote replication.

Viewing The Storage Array Profile

The Storage Array Profile displays the most detailed information about the components of a remote replication and the replication repository virtual disks. You can:

- View detailed information about individual and paired virtual disks in a remote replication
- View detailed information about the replication repository virtual disks in the remote replication and storage array
- Save the Storage Array Profile information as a text file

To view the Storage Array Profile:

1. In the AMW of the local or remote storage array, select the **Summary** tab.
2. In the **Monitor** area, click **View Storage Array Profile**.
   The **Storage Array Profile** window is displayed.
3. Under the **Storage** tab, select the **Virtual Disks** tab.
4. Select either the **Replications** tab or the **Repositories** tab.
   The Profile for the Storage array page is displayed.
5. Perform one of the following:
   - To return to the AMW without saving the information, click **Close**.
   - To save the information, click **Save As**, then go to Step 6.

6. In the **Section Selection** area of the **Save Profile** window, select each section you want to save.
7. To save the file, either:
   - Save the file in the default directory, or
   - Save the file in another directory by choosing it in the Look in drop-down list.
8. In **File name**, type a name for the file and click **Save**.
   The file is saved as a .txt file.

Viewing The Properties Pane

The Properties pane is a view-only display of the physical and logical characteristics of a single virtual disk in a replicated pair, or a single replication repository virtual disk.

To view the Properties pane:

1. In the AMW, select the **Storage & Copy Services** tab.
2. Select either the primary or secondary virtual disk of the replicated pair.
   The Properties pane shows properties for the selected virtual disk. Under **Replication** properties, the Replication status displays the synchronization status of the replicated pair. When the primary and secondary virtual disks are synchronizing, the Replication status shows a synchronizing icon. For more information, see Virtual Disk Status Icons.

Viewing Logical Elements Of The Secondary Virtual Disk

To view the logical elements of the secondary virtual disk of a remote replication:

1. In the AMW of the local storage array, select the **Storage & Copy Services** tab.
2. Right-click the secondary virtual disk of the remote replication.
3. Select **View Associated Logical Elements**.
   The **View Associated Logical Elements** pop-up is displayed, displaying these logical elements:
   - Primary and secondary virtual disks and their locations
   - Replication repository virtual disks and their locations

Viewing The Physical Components Or Logical Elements Of The Primary Virtual Disk

To view the physical components or logical elements of the primary virtual disk in a remote replication:

1. In the AMW of the storage array that contains t he primary virtual disk, select the **Storage & Copy Services** tab.
2. Right-click the primary virtual disk and do one of the following:
   - View the logical elements of the primary virtual disk by selecting **Select View → Associated Logical Elements.** The **View Associated Logical Elements** pop-up is displayed, showing visual representations of these elements: the primary and secondary virtual disk in the remote replication (and their locations) and the replication repository virtual disks in the storage array (and their locations).
   - View the physical components of the primary virtual disk by clicking **View Associated Physical Components** in the Properties pane. The **View Associated Physical Components** pop-up is displayed, showing a visual representation of the primary virtual disk in the remote replication.

Virtual Disk Status Icons

In the **Devices** tab on the EMW, the following icons depict the status of each virtual disk.

Icon	Status
	Replication repository virtual disk
	Primary virtual disk
	Secondary virtual disk
	Primary virtual disk, synchronization in progress
	Primary virtual disk, replication suspended
	Secondary virtual disk, replication suspended
	Primary virtual disk, unsynchronized
	Secondary virtual disk, unsynchronized with communications loss

Changing Write Mode And Consistency Group Membership

The write mode of a remote replication is selected when it is created. However, you can change the write mode in a remote replication at a later time, as well as change the secondary virtual disk's membership in a write consistency group. For more information about write modes and write consistency groups, see the online help topics.

To change Write Mode and Consistency Group Membership:

1. In the AMW the storage array containing the primary virtual disk, select the **Storage & Copy Services** tab.
2. In the object tree, right-click the primary virtual disk of the replicated pair.
3. Select **Change** → **Write Mode.**

   The **Change Write Mode** dialog is displayed. The Replicated pairs table reflects all replicated pairs in both the local and remote storage arrays.
4. Select one or more replicated pairs. To select all replicated pairs, click **Select All**.
5. Select either the **Synchronous** or **Asynchronous** write mode.
6. If you are you adding the secondary virtual disk of the replicated pair to a write consistency group, select **Yes** and then select **Add to consistency group**. If you are not adding the secondary virtual disk of the replicated pair to a write consistency group, go to Step 7.
7. Click OK.

   The **Change Write Mode** window is displayed.
8. In the **Change Write Mode** window, click **Yes**.

   The **Replication properties** section (the Properties panes in the AMW for the local storage array) displays the following information:

- Replication status is Synchronized.
- Write mode is either Synchronous or Asynchronous, depending on your selection above.
- The secondary virtual disk is either write consistent/not write consistent.
- Resynchronization method is either Manual or Automatic.

Resynchronizing Virtual Disks

There are two resynchronization methods for remote replications:

- Manual resynchronization — See Manually Resynchronizing Virtual Disks.
- Automatic resynchronization — See Automatically Resynchronizing Virtual Disks.

For more information about synchronization and resynchronization in remote replications, see:

- Normally Synchronized Virtual Disks.
- Unsynchronized Virtual Disks.
- Setting Synchronization Priority And Synchronization Method.
- Changing Synchronization Priority And Method.
- Resynchronizing Virtual Disks.

You may require to periodically test the communication between the primary virtual disk and the secondary virtual disk in a remote replication, especially after resynchronizing virtual disks. For information about testing the communication, go to Testing Communication Between The Primary And Secondary Virtual Disks.

Normally Synchronized Virtual Disks

In a normally synchronized remote replication, RAID controller module owners manage the transfer of data from the primary to the secondary virtual disk. In a normal remote replication, these events occur:

- The primary virtual disk receives a write request from a host.
- The RAID controller module owner on the storage array logs information about the write operation to a replication repository virtual disk in the storage array.
- The RAID controller module owner writes the data to the primary virtual disk.
- The RAID controller module owner starts a data transfer operation to the secondary virtual disk on the secondary storage array.

The communication between a primary and secondary virtual disks can be suspended or become unsynchronized. If communication between the primary virtual disk and the secondary virtual disk is disrupted:

- The status of the replicated pair changes to Unsynchronized.
- A Needs Attention status is displayed for the storage array.
- Data is written to the primary virtual disk.
- Write requests to the primary virtual disk are logged.
- The RAID controller module owner sends an I/O completion message to the host sending the write request. Although the host can continue to send write requests to the primary virtual disk, no data transfer takes place to the secondary virtual disk. Writes to the secondary virtual disk are suspended pending restoration of communications between the primary virtual disk and the secondary virtual disk.

When connectivity is restored between the primary virtual disk and the secondary virtual disk, the replicated pair is ready to be resynchronized.

NOTE: When the primary virtual disk and the secondary virtual disk are resynchronized, only data that has changed on the primary virtual disk after the break in communication is transferred to the secondary virtual disk.

CAUTION: You should be aware of a possible loss of data if communication is broken after resynchronization starts between the primary storage array and the secondary storage array. New data might mix with old data on the secondary virtual disk and render the data unusable in a disaster recovery situation.

Changing Synchronization Priority And Method

Synchronization priority defines how much processing time and resources are allocated to synchronizing the primary and secondary virtual disk of a remote replication relative to system performance. Increasing the synchronization priority of a remote replication may degrade system performance. Synchronization priorities can be reset at any time.

Synchronization priorities can affect these operations:

- Performing a copyback
- Performing a Virtual Disk Expansion (VDE)
- Reconstructing a virtual disk
- Initializing a virtual disk
- Changing the segment size of a virtual disk
- Defragmenting a disk group
- Adding free capacity to a disk group
- Changing the RAID level of a disk group

To change the synchronization priority and method for an existing remote replication:

1. In the AMW of the storage array containing the primary virtual disk of the replicated pair, right-click the **Storage & Copy Services** tab.
2. Select **Change → Synchronization Settings.**
   The **Change Synchronization Settings** dialog is displayed.
3. In the Replicated pairs table, select the primary virtual disk and the remote virtual disk you want to change. To select all virtual disks, click **Select All**.
4. On the **Select Synchronization Priority** slide bar, select the synchronization priority for the replicated pair.
5. Select either Manual or Automatic resynchronization.
   - Automatic — Resynchronization starts immediately after communication is restored between unsynchronized replicated virtual disks.
   - Manual — The replicated pair must be manually resynchronized each time communication is restored between unsynchronized replicated virtual disks.
6. Click **OK**.
   The **Change Synchronization Settings** dialog is displayed.
7. In the **Change Synchronization Settings** dialog, click **Yes**.
   The **Change Synchronization Priority - Progress** bar displays the progress of the resynchronization.
8. Click **OK**.

Unsynchronized Virtual Disks

Communication between a primary and secondary virtual disks can be either suspended, or become unsynchronized. If this communication between both virtual disks is disrupted:

- The status of the replicated pair changes to **Unsynchronized**.
- A **Needs Attention** status is displayed for the storage array.
- Data is written to the primary virtual disk.
- Write requests to the primary virtual disk are logged.
- The RAID controller module owner sends an I/O completion message to the host sending the write request. Although the host can continue to send write requests to the primary virtual disk, no data transfer takes place to the secondary virtual disk. Writes to the secondary virtual disk are suspended pending restoration of communications between the primary and secondary virtual disks.

When connectivity is restored between the primary and secondary virtual disks, the replicated pair is ready to be resynchronized.

NOTE: When the primary virtual disk and the secondary virtual disk are resynchronized, only data that has changed on the primary virtual disk after the break in communication is transferred to the secondary virtual disk.

CAUTION: Loss of data can occur if communication is broken after resynchronization starts between the primary storage array and the secondary storage array. In addition to resulting in an incomplete resynchronization that may render the secondary storage array unusable, an interrupted resynchronization could cause a loss of data integrity on both arrays.

Automatically Resynchronizing Virtual Disks

When automatic resynchronization is selected, the RAID controller module owning the primary virtual disk automatically starts resynchronizing the data on the remote replication pair immediately after communication is restored.

NOTE: When choosing automatic resynchronization, a possible loss of data can occur if a resynchronization is interrupted in progress. Since another resynchronization begins automatically after communication is restored between the primary virtual disk and the remote virtual disk, data integrity could be affected.

With automatic resynchronization set, you cannot add a secondary virtual disk to a write consistency group. Therefore, write consistency during the resynchronization process is not preserved. The write order is not consistent until the entire group achieves **Optimal** status.

Manually Resynchronizing Virtual Disks

When manual resynchronization is selected, you must manually resynchronize and resume the data transfer on a remote replication after communication is restored between the primary and remote virtual disks. Manual resynchronization is the recommended setting for all remote replications for these reasons:

- You determine when resynchronization starts, which allows you to better manage and/or mitigate impact on I/O performance.
- In a disaster recovery situation, manual resynchronization offers the best chance of retrieving valid data.
- When the secondary virtual disk is in a write consistency group, manual resynchronization preserves the write order.

Reversing Roles Between The Primary And Secondary Virtual Disks

Reversing the roles between the virtual disks promotes the secondary virtual disk to the role of primary virtual disk and demotes the primary virtual disk to the role of secondary virtual. If the primary virtual disk in a remote virtual disk replication fails, you can reverse the roles of the primary virtual disk and the secondary virtual disk to transfer the data back to the restored virtual disk.

CAUTION: A potential loss of data access can occur if you reverse roles between the secondary virtual disk and the primary virtual disk while a virtual disk copy is in progress. The role reversal may succeed, but the virtual disk copy fails and cannot be restarted.

- You cannot perform a virtual disk copy on a secondary virtual disk in a virtual disk remote replication. To create a virtual disk copy of a secondary virtual disk, you must reverse the roles of the secondary virtual disk and the primary virtual disk, then perform the virtual disk copy on the new primary virtual disk.
- While a virtual disk remote replication is synchronizing, you cannot perform a virtual disk copy on either the primary or secondary virtual disk.
- If you reverse roles between a secondary virtual disk with less capacity than the primary virtual disk, the role reversal succeeds but the usable capacity of the new primary virtual disk equals the total capacity of the secondary virtual disk.

Promoting The Secondary Virtual Disk Or Demoting The Primary Virtual Disk

To promote the secondary virtual disk to the role of primary virtual disk, or demote the primary virtual disk to the role of secondary virtual disk:

1. In the AMW of the storage array that contains the virtual disk you are changing, click the **Storage & Copy Services** tab.
2. Right-click the virtual disk you want to change.
   - If you are promoting the secondary virtual disk to the role of primary virtual disk, select **Change** → **Role** to Primary. The **Change to Primary** dialog is displayed. Clicking **Yes** reverses the roles of the primary and secondary virtual disks.
   - If you are demoting the primary virtual disk to the role of secondary virtual disk, select **Change Role** to Secondary. The **Change to Secondary** dialog is displayed. Clicking **Yes** reverses the roles of the primary and secondary virtual disks.

Suspending A Remote Replication (Legacy)

1. In the AMW of the storage array with the primary virtual disk, select the **Storage & Copy Services** tab.
2. Right-click the primary virtual disk of a replicated pair and select **Suspend Replication**.

   The **Suspend Replicated Pair** dialog is displayed. The Replicated pairs table shows all replicated pairs in the local and remote storage arrays.
3. Select one or more of the replicated pairs to suspend. To select all replicated pairs, click **Select All**.
4. Click **Suspend**.

   The `Suspend Replication Relationship - Confirmation` message is displayed.
5. In the **Confirmation** text box, type `Yes`, then click **OK**.

   The Suspend Replicated Pair - Progress bar indicates the progress of the suspension. The Properties pane in the AMW that contains the suspended primary virtual disk shows the Replication status as **Suspended**. The suspended icon is displayed next to both the primary virtual disk icon and the secondary virtual disk icon in the Logical pane in the AMW.

Resuming A Remote Replication (Legacy)

1. In the AMW of the storage array with the primary virtual disk, select the Logical tab.
2. In the Logical pane, right-click the primary virtual disk of the replicated pair, then select Resume Replication.

   The **Resume Replicated Pair** dialog is displayed. The Replicated pairs table shows all suspended replicated pairs in the local and remote storage arrays.

3. Select one or more replicated pairs. To select all replicated pairs, click **Select All**.
4. Click **Resume**.

   The **Resume Replication Pair - Confirmation** dialog is displayed.
5. In the **Confirmation** text box, type `Yes`.

   The remote virtual disk replication resumes.

   NOTE: The Properties panes in the AMW for the local storage array and the remote storage array show the Replication status as Synchronized for both the primary and secondary virtual disks.

About Resumed Remote Replications

When a remote virtual disk replication is suspended, data continues to write to the primary virtual disk, but is not replicated to the secondary virtual disk. Writes to the primary virtual disk are persistently logged to the replication repository virtual disks.

After communications are restored, data transfer between the primary virtual disk and the secondary virtual disk must be resynchronized.

Depending on which resynchronization method you specify, the following occurs:

- Automatic resynchronization — Data transfer automatically starts immediately after the virtual disks are resynchronized.
- Manual resynchronization — You must manually resume the remote virtual disk replication to restart the data transfer. A suspended remote virtual disk replication stays in a Suspended status until it is manually resumed.

After the remote virtual disk replication resumes, data is automatically written to the secondary virtual disk. Only the regions of the primary virtual disk that changed since the replicated pair was suspended are written to the secondary virtual disk.

CAUTION: Possible loss of data access: When you resume a remote virtual disk replication involving a primary or secondary virtual disk that is a member of a write consistency group, any other suspended remote replications for replicated pairs in the write consistency group also resumes.

NOTE: If the write mode is synchronous, you do not need to resynchronize the primary and secondary virtual disk after resuming a remote virtual disk replication.

Testing Communication Between The Primary And Secondary Virtual Disks

When a resynchronization is manual or during a disaster recovery scenario, you may want to test the communication between the primary virtual disk and secondary virtual disk in a remote replication. For more information about synchronization and resynchronization in remote virtual disk replications, go to these topics:

- Normally Synchronized Virtual Disks.
- Unsynchronized Virtual Disks.
- Setting Synchronization Priority And Synchronization Method.
- Changing Synchronization Priority And Method.
- Resynchronizing Virtual Disks.
- Automatically Resynchronizing Virtual Disks.

To test the communication between virtual disks in a remote replication:

1. In the AMW of either the primary or secondary virtual disk, select the **Storage & Copy Services** tab.
2. In the Logical pane, right-click the virtual disk.

3. Select **Test Replication Communication**.
   The `Replication Communication Test Progress` message is displayed.

Deleting A Virtual Disk From A Replicated Pair In A Storage Array

Follow these steps to delete either a primary virtual disk, a secondary virtual disk, or both virtual disks from a replicated pair in a storage array.

NOTE: Do not remove a replication relationship to back up a replicated virtual disk. To perform backups of either the primary virtual disk or the secondary virtual disk, suspend the remote virtual disk replication so that the replication relationship is not broken.

Deleting A Primary Virtual Disk

CAUTION: Depending on which premium features are enabled on the storage array, deleting a primary virtual disk might delete all associated virtual disks, resulting in a possible loss of data on those virtual disks.

NOTE: You cannot delete a primary virtual disk while it is synchronizing.

When a primary virtual disk is deleted from a remote virtual disk replication, the following occurs:

- The primary virtual disk is deleted from the storage array.

  NOTE: The virtual disk is permanently deleted from the storage array and all data on the primary virtual disk is permanently lost.
- The replication relationship breaks.
- The capacity of the deleted virtual disk becomes unconfigured free capacity in the storage array and is available for creation of new virtual disks.
- The secondary virtual disk becomes a regular, standard virtual disk and is can accept both reads and writes.

Deleting A Secondary Virtual Disk

CAUTION: Deleting a secondary virtual disk results in the permanent loss of the data on the secondary virtual disk.

NOTE: Depending on which premium features are enabled on the storage array, deleting a secondary virtual disk may delete all associated virtual disks, resulting in a possible loss of data on those virtual disks.

NOTE: You cannot delete a secondary virtual disk while it is synchronizing.

When a secondary virtual disk is deleted, the replication relationship is removed and the remote virtual disk replication is destroyed. For steps describing how to delete a secondary virtual disk, see Deleting A Virtual Disk From A Replicated Pair In A Storage Array.

Deleting A Primary Virtual Disk In A Replicated Pair From A Storage Array

1. Stop all I/O activity to the primary virtual disk and unmount any file systems on the primary virtual disk.
2. In the AMW of the storage array containing the primary virtual disk, select the **Storage & Copy Services** tab.
3. In the Logical pane, right-click the primary virtual disk, then select **Delete**.
   The **Delete Virtual Disks** dialog is displayed.
4. Select one or more virtual disks to delete, then click **Delete**.
   The **Confirm Delete Virtual Disk(s)** dialog is displayed.
5. In the **Confirmation** text box, type `Yes` and click **OK**.
   The **Delete Virtual Disks - Progress** bar is displayed.
6. When the deletion is complete, click **OK**.

The primary virtual disk is deleted from the storage array. The secondary virtual disk in the replicated pair now becomes a regular standard virtual disk in the storage array.

NOTE: The primary virtual disk is deleted from the storage array and all data on the virtual disk is permanently lost.

Deleting A Secondary Virtual Disk In A Replicated Pair From A Storage Array

1. Stop all I/O activity on the secondary virtual disk and unmount any file systems on the secondary virtual disk.
2. In the AMW of the storage array that contains the secondary virtual disk, select the **Storage & Copy Services** tab.
3. In the object tree, right-click the secondary virtual disk, then select **Delete**.
   The **Delete Virtual Disks** dialog is displayed.
4. Select one or more virtual disks to delete, then click **Delete**.
   The **Confirm Delete Virtual Disk(s)** dialog is displayed.
5. In the **Confirmation** text box, type `Yes` and click **OK**.
   The **Delete Virtual Disks - Progress** bar is displayed.
6. When the deletion is complete, click **OK**.
   The replication relationship is removed and the remote virtual disk replication is destroyed.

   NOTE: The secondary virtual disk is deleted from the storage array and all data on the virtual disk is permanently lost.

Removing A Remote Replication From A Storage Array

Removing a remote virtual disk replication from a storage array returns both the primary virtual disk and the secondary virtual disk to regular standard virtual disks. After the removal is complete, normal I/O operations continue on the former primary virtual disk; the former secondary virtual disk also becomes available for normal I/O operations. Both virtual disks become read-write enabled. Unless one of disks has been deleted, a replication relationship between the two virtual disks can be re-created.

CAUTION: Do not remove a replication relationship to back up a replicated virtual disk. To back up either the primary virtual disk or the secondary virtual disk, suspend the remote virtual disk replication so that the replication relationship is not broken.

NOTE: No data on either virtual disk is deleted.

To remove remote replication from a storage array:

1. In the AMW of the storage array containing the primary virtual disk, select the **Storage & Copy Services** tab.
2. In the Logical pane, right-click the primary virtual disk of a replicated pair, then select **Remove Replication Relationship**.
   The **Remove Replication Relationship** dialog is displayed. The **Replicated pairs** table shows all replicated pairs in both the local and remote storage array.
3. Do one of the following:
   - To remove a replication relationship, select one or more replicated pairs.
   - To select all replicated pairs, click **Select All**.
4. Click **Remove**.
   The **Remove Replication Relationship - Confirmation** dialog is displayed.
5. Click **Yes**.
   The **Remove Replicated Pair - Progress** bar shows the progress of the removal process.

Deactivating The Remote Replication (Legacy) On The Storage Array

Before you can permanently disable the Remote Replication premium feature you must deactivate it on the storage array. See Disabling The Remote Replication Premium Feature. Deactivating the Remote Replication premium feature on this storage array does not affect remote virtual disk replications or the Remote Replication premium features on other storage arrays. However, another storage array cannot use this storage array as a remote storage array for creating a separate remote virtual disk replication.

To deactivate remote replication in a storage array:

1. In the AMW, select **Copy Services** → **Remote Replication** → **Deactivate.**
   The `Deactivate Replication` confirmation message is displayed.
2. Click **Yes**.
   The Remote Replication premium feature is deactivated and the two replication repository virtual disks are deleted from the storage array.

Disabling The Remote Replication (Legacy) Premium Feature

Before disabling the Remote Replication premium feature, all remote virtual disk replications must have been removed on the storage array. See Disabling The Remote Replication Premium Feature.

NOTE: To re-enable the Remote Replication premium feature on a storage array after it has been disabled, you must either retrieve the premium feature key or obtain a new one from your Dell support representative.

To disable the remote replication premium feature:

1. In the AMW, select **Storage Array** → **Premium Features.**
   The **Premium Features and Feature Pack Information** screen is displayed.
2. Select **Remote Replication** in the **Premium Features** window and click **Disable**.
   A warning message is displayed.
3. If you are sure you want to disable the Remote Replication feature, click **Yes** to confirm.
   The Remote Replication premium feature is disabled.

17

Management Firmware Downloads

Downloading RAID Controller And NVSRAM Packages

A version number exists for each firmware file. The version number indicates whether the firmware is a major version or a minor version. You can use the Enterprise Management Window (EMW) to download and activate both the major firmware versions and the minor firmware versions. You can use the Array Management Window (AMW) to download and activate only the minor firmware versions.

NOTE: Firmware versions are of the format **aa.bb.cc.dd**. Where **aa** is the major firmware version. **bb.cc.dd** is the minor firmware version. Depending on which one changes, firmware can be updated from EMW and AMW or only EMW.

You can activate the files immediately or wait until a more convenient time. You may want to activate the firmware or NVSRAM files at a later time because of these reasons:

- Time of day — Activating the firmware and the NVSRAM can take a long time, so you can wait until I/O loads are lighter. The RAID controller modules are offline briefly to load the new firmware.
- Type of package — You may want to test the new firmware on one storage array before loading the files onto other storage arrays.

The ability to download both files and activate them later depends on the type of RAID controller module in the storage array.

NOTE: You can use the command line interface to download and activate the firmware to several storage arrays by using a script. See the online help topics for more information on the command line interface.

Downloading Both RAID Controller And NVSRAM Firmware

NOTE: I/O to the array can continue while you are upgrading RAID controller and NVSRAM firmware.

NOTE: It is recommended that the firmware and NVSRAM be upgraded during a maintenance period when the array is not being used for I/O.

NOTE: The RAID enclosure must contain at least two disk drives in order to update the firmware on the controller.

To download RAID controller and NVSRAM firmware in a single operation:

1. If you are using the EMW, go to step 9. If you are using the AMW, go to step 2.
2. In the AMW, select **Upgrade → RAID Controller Module Firmware → Upgrade.**
   The **Download RAID Controller Module Firmware** is displayed.

   NOTE: The **RAID Controller Module Firmware** area and the NVSRAM area list the current firmware and the current NVSRAM versions respectively.
3. To locate the directory in which the file to download resides, click **Select File** next to the **Selected RAID controller module firmware file** text box.
4. In the **File Selection** area, select the file to download.

By default, only the downloadable files that are compatible with the current storage array configuration are displayed.

When you select a file in the **File Selection** area of the dialog, applicable attributes (if any) of the file are displayed in the **File Information** area. The attributes indicate the version of the file.

5. If you want to download an NVSRAM file with the firmware:
   a) Select **Transfer NVSRAM file with RAID controller module firmware**.
   b) Click **Select File**.
6. To transfer the files to the RAID controller module without activating them, click **Transfer files but don't activate them (activate later)**.
7. Click **Transfer**.

   Keep these guidelines in mind:

   - If the **Transfer** button is inactive, ensure that you either select an NVSRAM file or clear the **Transfer NVSRAM file with RAID controller module firmware** check box.
   - If the file selected is not valid or is not compatible with the current storage array configuration, the **File Selection Error** dialog is displayed. Click **OK** to close it, and choose a compatible firmware or NVSRAM file.
8. In the **Confirm Download** dialog, click **Yes**.

   The download starts.
9. If you are using the EMW, perform one of these actions:

   - Select **Tools → Upgrade RAID Controller Module Firmware**.
   - Select the **Setup** tab, and click **Upgrade RAID Controller Module Firmware**.
10. In the **Storage array** pane, select the storage array for which you want to upgrade the RAID controller module firmware or the NVSRAM.

    You can select more than one storage array.

    NOTE: The Details pane shows the details of only one storage array at a time. If you select more than one storage array in the Storage Array pane, the details of the storage arrays are not shown in the Details pane.
11. Click **Firmware** in the **Download** area.

    If you select a storage array that cannot be upgraded, the **Firmware** button is disabled. The **Download Firmware** dialog is displayed. The current firmware version and the NVSRAM version of the selected storage arrays appear.

    NOTE: If you select the storage arrays with different RAID controller module types that cannot be updated with the same firmware or NVSRAM file and click Firmware, the Incompatible RAID Controller Modules dialog is displayed. Click OK to close the dialog and select the storage arrays with similar RAID controller module types.
12. To locate the directory in which the file to download resides, click **Browse** in the **Select files** area.

    The **Select File** dialog is displayed.
13. Select the file to download.
14. Click **OK**.
15. If you want to download the NVSRAM file with the RAID controller module firmware, select **Download NVSRAM file with firmware** in the **Select files** area.

    Any attributes of the firmware file are displayed in the Firmware file information area. The attributes indicate the version of the firmware file.

    Any attributes of the NVSRAM file are displayed in the NVSRAM file information area. The attributes indicate the version of the NVSRAM file.
16. If you want to download the file and activate the firmware and NVSRAM later, select the **Transfer files but don't activate them (activate later)** check box.

NOTE: If any of the selected storage arrays do not support downloading the files and activating the firmware or NVSRAM later, the **Transfer files but don't activate them (activate later)** check box is disabled.

17. Click **OK**.

    The **Confirm Download** dialog is displayed.

18. Click **Yes**.

    The download starts and a progress indicator is displayed in the Status column of the **Upgrade RAID Controller Module Firmware** window.

Downloading Only NVSRAM Firmware

Use the command line interface (CLI) to download and activate NVSRAM to several storage arrays. For more information, see the online help.

To download only NVSRAM firmware:

1. To download the NVSRAM firmware from:
   - EMW — Go to step 7.
   - AMW — Go to step 2.
2. In the AMW, select **Upgrade → RAID Controller Module NVSRAM**

   or

   select the **Support** tab, and click **Download Firmware**. In **Select download task**, select **Download RAID controller module NVSRAM** and click **OK**. An error message is displayed. Click **OK** to close it, and select a compatible file.
3. To locate the directory in which the file to download resides, click **Select File**.
4. Select the file to download in the File selection area, and click **OK**.

   By default, only downloadable files that are compatible with the current storage array configuration are displayed.

   When you select a file in the File selection area, applicable attributes (if any) of the file appear in the NVSRAM File information area. The attributes indicate the version of the NVSRAM file.
5. Click **Transfer**.

   NOTE: If the file selected is not valid or is not compatible with the current storage array configuration, the **File Selection Error** dialog is displayed. Click **OK** to close it, and choose a compatible NVSRAM file.
6. Click Yes in the **Confirm Download** dialog.

   The download starts.
7. Perform one of these actions:
   - Select **Tools → Upgrade RAID Controller Module Firmware**.
   - Select the **Setup** tab, and click **Upgrade RAID Controller Module Firmware**.

   The **Upgrade RAID Controller Module Firmware** window is displayed.

   The Storage array pane lists the storage arrays. The Details pane shows the details of the storage array that is selected in the Storage array pane.
8. In the Storage array pane, select the storage array for which you want to download the NVSRAM firmware.

   You can select more than one storage array.

   NOTE: The Details pane shows the details of only one storage array at a time. If you select more than one storage array in the Storage array pane, the details of the storage arrays are not shown in the Details pane.
9. Click **NVSRAM** in the **Download** area.

   NOTE: If you select a storage array that cannot be upgraded, the NVSRAM button is disabled.

The **Download NVSRAM** dialog is displayed. The current firmware version and the NVSRAM version of the selected storage arrays is displayed.

NOTE: If you select the storage arrays with different RAID controller module types that cannot be updated with the same NVSRAM file and click NVSRAM, the **Incompatible RAID Controller Modules** dialog is displayed. Click OK to close the dialog and select the storage arrays with similar RAID controller module types.

10. To locate the directory in which the NVSRAM file to download resides, click **Browse** in the **Select file** area.
    The **Select File** dialog is displayed.
11. Select the file to download.
12. Click **OK**.
    Attributes of the NVSRAM file are displayed in the NVSRAM file information area. The attributes indicate the version of the NVSRAM file.
13. Click **OK**.
    The **Confirm Download** dialog is displayed.
14. Click **Yes**.
    The download starts and a progress indicator is displayed in the Status column of the **Upgrade RAID Controller Module Firmware** window.

Downloading Physical Disk Firmware

CAUTION: When updating physical disk firmware, you should stop all I/O activity to the array to prevent data loss.

The physical disk firmware controls various features of the physical disk. The disk array controller (DAC) uses this type of firmware. Physical disk firmware stores information about the system configuration on an area of the physical disk called DACstore. DACstore and the physical disk firmware enable easier reconfiguration and migration of the physical disks. The physical disk firmware performs these functions:

- The physical disk firmware records the location of the physical disk in an expansion enclosure. If you take a physical disk out of an expansion enclosure, you must insert it back into the same physical disk slot, or the physical disk firmware cannot communicate with the RAID controller module or other storage array components.
- RAID configuration information is stored in the physical disk firmware and is used to communicate with other RAID components.

CAUTION: Risk of application errors—Downloading the firmware could cause application errors.

Keep these important guidelines in mind when you download firmware to avoid the risk of application errors:

- Downloading firmware incorrectly could result in damage to the physical disks or loss of data. Perform downloads only under the guidance of your Technical Support representative.
- Stop all I/O to the storage array before the download.
- Make sure that the firmware that you download to the physical disks are compatible with the physical disks that you select.
- Do not make any configuration changes to the storage array while downloading the firmware.

NOTE: Downloads can take several minutes to complete. During a download, the **Download Physical Disk - Progress** dialog is displayed. Do not attempt another operation when the **Download Physical Disk - Progress** dialog is displayed.

To download Physical Disk Firmware:

1. From the AMW, select **Upgrade → Physical Disk Firmware.**

The **Download Physical Disk Firmware - Introduction** window is displayed.

2. Click **Next**.
   The **Download Physical Disk Firmware - Add Package** window is displayed.
3. In the **Selected Packages** area, click **Add**. Navigate to the location of the packages and click **OK**.
   The selected package is added to the **Packages to be transferred** area.
4. Click **Next**.
   The **Download Physical Disk Firmware - Select Physical Disks** window is displayed.
5. In the **Compatible Physical Disks** tab, select the appropriate physical disks or **Select all** the physical disks.
   The **Confirm Download** dialog is displayed.
6. Type `yes` and click **OK**.
   The **Download Physical Disk Firmware - Progress** window displays the progress of physical disk firmware download.
7. After the firmware download is complete, click **Close**.

For more information, see the online help topics.

Downloading MD3060e Series Expansion Module EMM Firmware

NOTE: Do not make any configuration changes to the storage array while you are downloading the expansion enclosure EMM firmware. Doing so could cause the firmware download to fail, damage the storage array, or cause loss of data accessibility.

NOTE: Due to a limitation with Linux, expansion enclosure EMM firmware updates must be performed using out-of-band management only. Failure to do so may result in the host server becoming unresponsive, and it may require a reboot.

You can transfer a downloadable firmware file to the expansion enclosure EMM in the expansion enclosures attached to the storage array.

CAUTION: Risk of possible loss of data or risk of damage to the storage array—Downloading the expansion enclosure EMM firmware incorrectly could result in loss of data or damage to the storage array. Perform downloads only under the guidance of your Technical Support representative.

CAUTION: Risk of making expansion enclosure EMM unusable—Do not make any configuration changes to the storage array while downloading expansion enclosure EMM firmware. Doing so could cause the firmware download to fail and make the selected expansion enclosure unusable.

1. In the AMW, select**Upgrade → EMM Firmware**.
   The Download Environmental (EMM) Card Firmware dialog is displayed.
2. In the **Select enclosures** area, either select each expansion enclosure to which you want to download firmware, or select the Select All to select all of the expansion enclosures in the storage array.
   Each selected expansion enclosure must have the same product ID.
3. Click **Select File** to select the EMM firmware file.
   The **Select Environmental (EMM) Card Firmware File** dialog is displayed.
4. Select the file to download, and click **OK**.
5. Click **Start**.
6. Click **Yes** to continue with the firmware download.

   NOTE: If you click **Stop** while a firmware download is in progress, the download-in-progress finishes before the operation stops. The status for the remaining expansion enclosures changes to **Canceled**.

7. Monitor the progress and completion status of the download to the expansion enclosures. The progress and status of each expansion enclosure that is participating in the download is displayed in the Status column of the Select enclosures table.

   NOTE: Each firmware download can take several minutes to complete.

8. Perform one of these actions depending on whether the download succeeded:

   – The download succeeded — The statuses of all the expansion enclosures show Complete. You can close the **Download environmental (EMM) Card Firmware** dialog by clicking **Close**. The expansion enclosure EMM cards are now operating with the new firmware.

   – The download failed — The status of one expansion enclosure shows Failed, and the remainder of the expansion enclosures show Canceled. Make sure that the new firmware file is compatible before attempting another firmware download.

Self-Monitoring Analysis And Reporting Technology (SMART)

Self-Monitoring Analysis and Reporting Technology (SMART) monitors the internal performance of all physical disk components to detect faults indicating the potential for physical disk failure. SMART uses this information to report whether failure is imminent so that a physical disk can be replaced before failure occurs. The RAID controller monitors all attached drives and notifies users when a predicted failure is reported by a physical disk.

Media Errors And Unreadable Sectors

If the RAID controller detects a media error while accessing data from a physical disk that is a member of a disk group with a redundant RAID level (RAID 1, RAID 5 or RAID 10), the controller tries to recover the data from peer disks in the disk group and uses recovered data to correct the error. If the controller encounters an error while accessing a peer disk, it is unable to recover the data and affected sectors are added to the unreadable sector log maintained by the controller. Other conditions under which sectors are added to the unreadable sector log include:

- A media error is encountered when trying to access a physical disk that is a member of a nonredundant disk group (RAID 0 or degraded RAID 1, RAID 5 or RAID 10).
- An error is encountered on source disks during rebuild.

NOTE: Data on an unreadable sector is no longer accessible.

18

Firmware Inventory

A storage array is made up of many components, which may include RAID controller modules, physical disks, and enclosure management modules (EMMs). Each of these components contains firmware. Some versions of the firmware are dependent on other versions of firmware. To capture information about all of the firmware versions in the storage array, view the firmware inventory.

If the firmware inventory does not contain information for a particular storage array, the firmware inventory service is not available on that storage array.

You can also save the firmware inventory to a text file. You can then send the file to your Technical Support representative for analysis. Your Technical Support representative can detect any firmware mismatches.

Viewing The Firmware Inventory

1. Perform one of these actions based on whether you want to view the firmware information for one storage array or all storage arrays:
   - Single storage array — From the AMW, select **Summary** → **View Firmware Inventory.**
   - All storage arrays — From the EMW, select **Tools** → **Firmware Inventory.**
2. To save the firmware inventory to a text file, click **Save As**.

   NOTE: The suffix *.txt is added to the file name automatically if you do not specify a suffix for the file name.
3. In **File name** dialog box, enter a name for the file to be saved. You may also specify another physical disk and directory if you want to save the file in a location other than the default.
4. Click **Save**.

   An ASCII text file that contains the firmware inventory is saved to the designated directory.

19

System Interfaces

Virtual Disk Service

The Microsoft Virtual Disk Service (VDS) is a component of the Windows operating system. The VDS component utilizes third-party vendor specific software modules, known as providers, to access and configure third-party storage resources, such as MD Series storage arrays. The VDS component exposes a set of application programming interfaces (APIs) that provides a single interface for managing disks and other storage hardware. The MD Series VDS Provider enables Windows tools, including the Disk Manager, to access and configure storage array virtual disks.

The VDS Provider for the MD Series storage arrays is available on the MD Series resource DVD. For more information on VDS, see **microsoft.com**.

Volume Shadow-Copy Service

The Microsoft Volume Shadow-copy Service (VSS) is a component of the Microsoft Windows operating system. The VSS component utilizes third-party vendor specific software modules, known as providers, to access and utilize snapshot and disk copy functionality provided by third-party storage resources, such as MD Series storage arrays. The combination of the VSS component and the VSS Provider, included on the MD Series Resource media, enables the MD Series storage arrays to be utilized by third-party and Windows backup and snapshot applications.

NOTE: Virtual disks used as source virtual disks for VSS snapshots must not have names longer than 16 characters.

The VSS hardware provider uses the source virtual disk name as a prefix for the snapshot and repository virtual disk names. The resulting snapshot and repository names are too long if the source virtual disk name exceeds 16 characters.

VSS attaches to the service and uses it to coordinate the creation of snapshot virtual disks on the storage array. VSS-initiated snapshot virtual disks can be triggered through backup tools, known as requestors. The VSS Provider Configuration Tool makes available the following configuration options:

- Snapshot Repository Virtual Disk Properties — This section contains a drop-down list for the RAID level and a field for entering source virtual disk capacity percentage for snapshot repositories.
- Snapshot Repository Virtual Disk Location — This section contains a list of preferences for the location of the snapshot repository virtual disk. These preferences are honored whenever conditions permit.

The Microsoft VSS installer service for storage provisioning is available on the MD Series resource media in the **\windows\VDS_VSS** directory.

NOTE: When registering VSS during your Windows setup, the registration graphical user interface (GUI) prompts you to provide the name of your array because settings in the GUI are array-specific, not host-specific.

Storage Management VSS Hardware Provider tips:

- The number of snapshot virtual disks that can be created using a single snapshot set varies with the I/O load on the RAID controller modules. Under little or no I/O load, the number of virtual disks in a snapshot set must be limited to eight. Under high I/O loads, the limit must be three.
- The snapshot virtual disks created in the MD Storage Manager are differential snapshots. Plex snapshots are not supported.
- Virtual disks to be used as source virtual disks for VSS snapshots must not have names longer than 16 characters. The VSS hardware provider uses the base virtual disk name as a prefix for the snapshot and

repository virtual disk names. The resulting snapshot and repository names are too long if the source virtual disk name exceeds 16 characters.

NOTE: A volume is another term for virtual disk.

For more information on VDS and VSS, see **microsoft.com**.

20

Storage Array Software

Start-Up Routine

Look and listen during the array's start-up routine for the indications described in the table below. For a description of the front- and back-panel indicators, see About Your Storage Array.

Look/Listen for	Action
Alert messages	See your storage management documentation.
An unfamiliar constant scraping or grinding sound when you access a physical disk	See Getting Help.

Device Health Conditions

When you open the Enterprise Management Window (EMW), the Dell PowerVault Modular Disk Storage Manager (MD Storage Manager) establishes communication with each managed storage array and determines the current storage array status. The current status is represented by icons next to the managed storage array.

The status icons shown in the Tree view in the EMW represent a summary status for each storage array. If a storage array has a status of Needs Attention or a status of Fixing, determine the condition that is causing this status before attempting any management actions. You can determine the condition causing the Needs Attention status or the Fixing status by selecting the storage array and launching its Array Management Window (AMW).

After the AMW opens, select the **Hardware** tab to see the components in the storage array. A component that has a problem is indicated by a status icon.

The status icons indicate the status of the components that comprise the storage array. Also, the Recovery Guru option provides a detailed explanation of the conditions and the applicable steps to remedy any Needs Attention status. For more information, see Recovery Guru.

For the status of a storage array, the icons shown in the following table are used in the Tree view, the Table view, and both the EMW Status Bar and the AMW Status Bar.

Table 9. Status Icons and Description

Status	Icon	Description
Optimal		Each component in the managed storage array is in the desired working condition.
Needs Attention		There is a problem with the managed storage array that requires your intervention to correct it.

Status	Icon	Description
Unresponsive		The storage management station cannot communicate with the storage array or one RAID controller module or both RAID controller modules in the storage array.
Fixing Status		A Needs Attention status has been corrected and the managed storage array is currently transitioning to an Optimal state.
Unsupported		The node is currently not supported by this version of MD Storage Manager.
Software Unsupported		The storage array is running a level of software that is no longer supported by the MD Storage Manager.

In the Table view, every managed storage array is listed once, regardless of the number of attachments it has in the Tree view. After the storage array has been contacted by the MD Storage Manager, an icon representing its hardware status is displayed. Hardware status can be Optimal, Needs Attention, or Fixing. If, however, all of the network management connections from the storage management station to the storage array shown in the Tree view are Unresponsive, the storage array status is represented as Unresponsive.

In the EMW Status Bar and the AMW Status Bar, the icons also have these behaviors:

- Hold the mouse over the icon in the EMW Status Bar and the AMW Status Bar to show a tooltip with a brief description of the status.
- The icons for the Needs Attention status and Unresponsive status are displayed in the EMW Status Bar and the AMW Status Bar if there are discovered storage arrays with either condition.

The EMW Tree view has additional status icons that are shown in the following table.

Table 10. Additional Status Icons and Description

Status	Icon	Description
Unsupported Alerts with a Needs Upgrade Status		Setting an alert on a storage array with a Needs Upgrade status is not supported. In this case, the storage array shows both a Needs Upgrade status and an Unsupported Alerts icon in the Tree view. The Unsupported Alerts icon indicates that the storage array cannot be monitored.
Alert Set		You can set alerts at any of the nodes in the Tree view. Setting an alert at a parent node level, such as at a host level, sets alert for any child nodes. If you set an alert at a parent node level and any of the in-band storage array child nodes have a Needs Upgrade status, the Alert Disables status icon is displayed next to the parent node in the tree view.

Status	Icon	Description
Setting an Alert at the Parent Node Level		You can set alerts at any of the nodes in the Tree view. Setting an alert at a parent node level, such as at a host level, sets alert for any child nodes. If you set an alert at a parent node level and any of the in-band storage array child nodes have a Needs Upgrade status, the Alert Disables status icon appears next to the parent node in the tree view.
Adding a Storage Array		The Contacting Storage Array icon is shown in the Tree view and Table view until the current status of each managed storage array is known. The Contacting Storage Array icon is shown in the EMW Status Bar and the AMW Status Bar and the tooltip shows Contacting Storage arrays. As each storage array is contacted, its current status is obtained and shown in the Tree view and Table view. The applicable statuses are the Optimal, Needs Attention, Fixing, or Unresponsive.
Adding a Storage Array OK		No problems were encountered while adding the storage array. The MD Storage Manager continues to check for any status change events.
Adding a Storage Array Error		Displayed only when an error occurs.

In the Tree view, icons can appear in a string to convey more information. For example, the following string means that the storage array is optimal, an alert is set for the storage array, and firmware is available for download.

NOTE: The MD Storage Manager may take a few minutes to update a status change to Unresponsive or from Unresponsive. A status change from or to Unresponsive depends on the network link to the storage array. All other status change updates faster.

Trace Buffers

Trace information can be saved to a compressed file. The firmware uses the trace buffers to record processing activity, including exception conditions, that may be useful for debugging. Trace information is stored in the current buffer and can be moved to the flushed buffer after being retrieved. Because each RAID controller module has its own buffer, there may be more than one flushed buffer. The trace buffers can be retrieved without interrupting the operation of the storage array and with minimal effect on performance.

NOTE: Use this option only under the guidance of a Technical Support representative.

A zip-compressed archive file is stored at the location you specify on the host. The archive contains trace files from one or both of the RAID controller modules in the storage array along with a descriptor file named **trace_description.xml**. Each trace file includes a header that identifies the file format to the analysis software used by the Technical Support representative. The descriptor file contains:

- The WWN for the storage array.
- The serial number of each RAID controller module.
- A time stamp.
- The version number for the RAID controller module firmware.
- The version number for the management application programming interface (API).
- The model ID for the RAID controller module board.
- The collection status for each RAID controller module. If the status is Failed, the reason for failure is noted, and there is no trace file for the failed RAID controller module.

Retrieving Trace Buffers

To retrieve the trace buffers:

1. From the AMW, select **Monitor → Health → Retrieve Trace Buffers.**

   The **Retrieve Trace Buffers** dialog is displayed.
2. Select either **RAID controller module 0**, **RAID controller module 1**, or both.

   If the RAID controller module status message to the right of a check box indicates that the RAID controller module is offline, the check box is disabled.
3. From the **Trace buffers** list, select the relevant option.
4. To move the buffer, select **Move current trace buffer to the flushed buffer after retrieval**.

   NOTE: Move current trace buffer to the flushed buffer after retrieval is not available if the **Flushed buffer** option is selected in step 3.
5. Enter a name for the physical disk data filename in **Specify filename** or click **Browse** to navigate to a previously saved file to overwrite an existing file.
6. Click **Start**.

   The trace buffer information is archived to the file specified.
7. After the retrieval process is completed:
   - To retrieve trace buffers again using different parameters, repeat step 2 through step 6.
   - To close the dialog, click **Close**.

Collecting Physical Disk Data

You can use the **Collect Physical Disk Data** option to collect log sense data from all the physical disks on your storage array. Log sense data consists of statistical information that is maintained by each of the physical disks in your storage array. Your Technical Support representative can use this information to analyze the performance of your physical disks and for troubleshooting problems that may exist.

WARNING: Use this option only under the guidance of your Technical Support representative.

To collect physical disk data:

1. In the AMW, perform one of these actions:
   - To collect data from all of the physical disks in the storage array, select **Monitor → Health → Collect Physical Disk Data → All Physical Disks.**

– To collect data from a single physical disk that is selected in the **Hardware** tab, select **Monitor → Health → Collect Physical Disk Data → Selected Physical Disks.**

The **Collect Physical Disk Data** window is displayed.

2. Enter a name for the physical disk data filename in **Specify filename** or click **Browse** to navigate to a previously saved file to overwrite an existing file.

   The suffix *.bin is added to the file automatically if you do not specify a suffix for the file.

3. Click **Start**.

   The physical disk data collection is completed and saved at the location that you entered.

4. Click **OK**.

Creating A Support Data Collection Schedule

To creating a support data collection schedule:

1. From the EMW, select **Tools → Collect Support Data → Create/Edit Schedule.**

   The **Schedule Support Data Collection** dialog is displayed.

2. In the **Storage arrays** table, select one or more storage arrays for which you want to create a schedule.
3. Click the **Create/Edit** button.

   The **Create/Edit Schedule** dialog is displayed.

4. Select your desired settings and click **OK**.

   The **Schedule Support Data Collection** dialog is displayed. The **Storage arrays** table is updated with the schedule changes you made.

5. Select where you want to save the collected support data files:

   – To use the default location, select **Use default location**.
   – To choose another location, select **Use alternate location**, then click the **Browse** button to select the desired directory.

   NOTE: The filename cannot be modified.

6. Click **OK**.

Suspending Or Resuming A Support Data Collection Schedule

Suspending a support data collection schedule temporarily disables the scheduled operation. When you suspend a support data collection schedule, the schedule's timer continues to run, but the scheduled support data collections do not occur. Suspending a schedule does not affect the automatic collection of support data during major event log (MEL) events.

Resuming a schedule restarts the collection of support data on a scheduled basis. You can resume a suspended schedule at any time.

1. From the EMW, select **Tools → Collect Support Data → Create/Edit Schedule.**

   The **Schedule Support Data Collection** dialog is displayed.

2. In the **Storage arrays** table, select one or more storage arrays.
3. Perform one of the following actions:

   – To suspend a support data collection schedule, click **Suspend**, then click **Yes**.
   – To restart a support data collection schedule, click **Resume**, then click **OK**.

4. Click **OK**.

Removing A Support Data Collection Schedule

To remove a support data collection schedule:

1. From the EMW, select **Tools → Collect Support Data → Create/Edit Schedule.**
   The **Schedule Support Data Collection** dialog is displayed.
2. In the **Storage arrays table**, select one or more storage arrays.
3. Click **Remove**.
4. Review the information, then click **Yes**.
   The **Schedule Support Data Collection** dialog is displayed.
5. Click **OK**.

Event Log

You can use the Event Log Viewer to view a detailed list of events that occur in a storage array. The event log is stored on reserved areas on the storage array disks. It records configuration events and storage array component failures. The event log stores approximately 8000 events before it replaces an event with a new event. If you want to keep the events, you may save them, and clear them from the event log.

The MD Storage Manager records the following events:

- Critical events — Errors occurring on the storage array that needs to be addressed immediately. Loss of data access may occur if the error is not immediately corrected.
- Warning events — Errors occurring on the storage array resulting in degraded performance or reduced ability to recover from additional errors. Access to data has not been lost, but the must be corrected to prevent possible loss of data access in the event of an additional error.
- Informational events — Events occurring on the storage array that do not impact normal operations. This event is reporting a change in configuration or other information useful in evaluating the performance of the storage array.
- Debug events — Events occurring on the storage array that provides information useful in determining steps or states that led to the error. This information may be useful to your Technical Support representative in helping determine error causes.

The event log window has the following event views:

- Summary view — Shows an event summary in a tabular format.
- Detail view — Shows details about a selected event.

Viewing The Event Log

WARNING: Use this option only under the guidance of your Technical Support representative.

To view the event log:

1. In the AMW, select **Monitor → Reports → Event Log.**
   The Event Log is displayed. By default, the summary view is displayed.
2. To view the details of each selected log entry, select **View details**.
   A detail pane is added to the event log that contains detailed information about the log item. You can view the details about a single log entry at a time.
3. To save the event log, click **Save As**.

The **Save Events** dialog is displayed, navigate to the relevant folder, enter the relevant **file name**, and click **Save**.

4. To erase all log entries from the event log, click **Clear All**.
5. To exit the event log, click **Close.**

   For more information, see the online help topics.

Recovery Guru

The Recovery Guru is a component of MD Storage Manager that diagnoses critical events on the storage array and recommends step-by-step recovery procedures for problem resolution.

In the AMW, to display the Recovery Guru, perform one of these actions:

- Select **Monitor** → **Health** → **View Health (Recovery Guru).**
- On the **Summary** tab, click the **Storage Array Needs Attention** link.

You can detect a problem using the following indicators:

- Non-Optimal status icons
- Alert notification messages that are sent to the appropriate destinations
- Hardware indicator lights

The status icons return to Optimal status as problems are resolved.

Storage Array Profile

The storage array profile provides a description of all of the components and properties of the storage array. The storage array profile also provides the option to save the storage array profile information to a text file. You may want to use the storage array profile as an aid during recovery or as an overview of the current configuration of the storage array. Create a new copy of the storage array profile if your configuration changes.

1. To open the storage array profile, in the AMW, perform one of the following actions:
   - Select **Monitor** → **Reports** → **Storage Array Profile.**
   - Select the **Summary** tab, and click **View Storage Array Profile** in the **Monitor** area.

   The **Storage Array Profile** dialog is displayed. The **Storage Array Profile** dialog contains several tabs, and the title of each tab corresponds to the subject of the information contained.
2. Perform one of these actions in the **Storage Array Profile** dialog:
   - View detailed information — Go to step 3.
   - Search the storage array profile — Go to step 4.
   - Save the storage array profile — Go to step 5.
   - Close the storage array profile — Go to step 6.
3. Select one of the tabs, and use the horizontal scroll bar and the vertical scroll bar to view the storage array profile information.

   NOTE: You can use the other steps in this procedure to search the storage array profile, to save the storage array profile, or to close the storage array profile.
4. To search the storage array profile, perform these steps:

   a) Click .

   b) Type the term that you want to search for in the **Find** text box.

If the term is located on the current tab, the term is highlighted in the storage array profile information.

NOTE: The search is limited to the current tab. If you want to search for the term in other tabs, select the tab and click the **Find** button again.

c) Click the **Find** button again to search for additional occurrences of the term.

5. To save the storage array profile, perform these steps:
   a) Click **Save As**.
   b) To save all sections of the storage array profile, select **All sections**.
   c) To save information from particular sections of the storage array profile, select the **Select sections**, and select the check boxes corresponding to the sections that you want to save.
   d) Select an appropriate directory.
   e) In **File Name**, type the file name of your choice. To associate the file with a particular software application that opens it, specify a file extension, such as .txt.

      NOTE: The file is saved as ASCII text.

   f) Click **Save**.
6. To exit the storage array profile, click **Close**.

Viewing The Physical Associations

You can use the **Associated Physical Components** option to view the physical components that are associated with source virtual disks, snapshot virtual disks, snapshot repository virtual disks, disk groups, unconfigured capacity, and free capacity in a storage array.

To view the physical associations:

1. In the AMW, select a node in the **Storage & Copy Services** tab or in the object tree of the **Host Mappings** tab.
2. Click **View Associated Physical Components**. Alternatively, if the selected node is a virtual disk, right-click the node to open a pop-up menu and select **View → Associated Physical Components.** If the selected node is a disk group, unconfigured capacity, or free capacity, right-click the node to open a pop-up menu and select **View Associated Physical Components**.

   The **View Associated Physical Components** dialog is displayed with blue dots next to the physical components that are associated with the selected node.
3. To close the **View Associated Physical Components** dialog, click **Close**.

Recovering From An Unresponsive Storage Array Condition

A storage array can have an Unresponsive status for several reasons. Use the procedure in this topic to determine a possible cause and solution. The MD Storage Manager can take up to five minutes to detect that a storage array has become unresponsive or becomes responsive again. Before completing this procedure, make sure that you wait for some time before you decide that the storage array is still unresponsive.

To recover from an unresponsive storage array:

1. Check the Tree View in the EMW to see if all storage arrays are unresponsive.
2. If any storage arrays are unresponsive, check the storage management station network connection to make sure that it can reach the network.
3. Ensure that the RAID controller modules are installed and that there is power to the storage array.
4. If there a problem with the storage array, correct the problem.
5. Perform one of these actions, depending on how your storage array is managed:

   – Out-of-band managed storage array — Go to step 6.
   – In-band managed storage array — Go to step 12.

6. For an out-of-band managed storage array, ensure that the RAID controller modules are network accessible by using the ping command to make sure that the RAID controller module can be reached. Type one of these commands, and press <Enter>.
   - `ping <host-name>`
   - `ping <RAID controller module-IP-addres>`
7. If the verification is successful, see step 8, if not, see step 9.
8. Remove the storage array with the Unresponsive status from the EMW, and select **Add Storage Array** to add the storage array again.
9. If the storage array does not return to Optimal status, check the Ethernet cables to make sure that there is no visible damage and that they are securely connected.
10. Make sure the appropriate network configuration tasks have been performed. For example, make sure that IP addresses have been assigned to each RAID controller module.
11. If there is a cable or network accessibility problem, see step 20, if not step 12.
12. For an in-band managed storage array, make sure that the host is network accessible by using the `ping` command to verify that the host can be reached. Type one of these commands, and press <Enter>.
    - `ping <host-name>`
    - `ping <RAID controller module-IP-addres>`
13. If the verification is successful, see step 14, if not, step 15.
14. Remove the host with the Unresponsive status from the EMW, and select **Add Storage Array** to add the host again.
15. If the host does not return to Optimal status, go to step 16.
16. Ensure that the host is turned on and operational and that the host adapters have been installed.
17. Check all external cables and switches or hubs to make sure that no visible damage exists and that they are securely connected.
18. Make sure the Host Context Agent software is installed and running.

    If you started the host system before you were connected to the RAID controller module in the storage array, the Host Context Agent software will not be able to detect the RAID controller modules. If this is the case, make sure that the connections are secure, and restart the Host Context Agent software.
19. If you have recently replaced or added the RAID controller module, restart the Host Context Agent software so that the new RAID controller module is recognized.
20. If the problem still exists, make the appropriate host modifications, check with other administrators to see if a firmware upgrade was performed on the RAID controller module from another storage management station.

    If a firmware upgrade was performed, the EMW on your management station may not be able to locate the new AMW software needed to manage the storage array with the new version of the firmware.
21. If the problem persists contact your Technical Support representative.
22. Determine if there is an excessive amount of network traffic to one or more RAID controller modules.

    This problem is self-correcting because the EMW software periodically retries to establish communication with the RAID controller modules in the storage array. If the storage array was unresponsive and a subsequent attempt to connect to the storage array succeeds, the storage array becomes responsive.

    For an out-of-band managed storage array, determine if management operations are taking place on the storage array from other storage management stations. A RAID controller module-determined limit exists to the number of Transmission Control Protocol/Internet Protocol (TCP/IP) connections that can be made to the RAID controller module before it stops responding to subsequent connection attempts. The type of management operations being performed and the number of management sessions taking place together determine the number of TCP/IP connections made to a RAID controller module. This problem is self-correcting because, after some TCP/IP connections terminate, the RAID controller module then becomes responsive to other connection attempts.
23. If the storage array is still unresponsive, a problem may exist with the RAID controller modules. Contact your Technical Support representative.

Locating A Physical Disk

You can physically locate and identify one or more of the physical disks in an expansion enclosure by activating physical disk LEDs.

To locate the physical disk:

1. Select the **Hardware** tab.
2. Select the physical disks that you want to locate.
3. Select **Hardware → Blink → Physical Disk.**
   The LEDs on the selected physical disks blink.
4. When you have located the physical disks, click **OK**.
   The LEDs stop blinking. If any other blink operations (Blink Disk Group, Blink Storage Array, Blink Physical Disk Ports, or Blink Expansion Enclosure) are currently being invoked from another storage management station, these LEDs also stop blinking.
5. In the rare case that the LEDs on the physical disks do not stop blinking, in the AMW, select **Hardware → Blink → Stop All Indications.**
   If the LEDs successfully stop blinking, a confirmation message is displayed.
6. Click **OK**.

Locating An Expansion Enclosure

You can use the **Blink** option to physically locate and identify an expansion enclosure in the storage array.

The LED activation varies according to the type of expansion enclosure that you have.

- If you have an expansion enclosure with a white LED, the Blink Expansion Enclosure operation causes the white LED on the expansion enclosure to come on. The LED does not blink.
- If you have any other types of expansion enclosures, this operation causes the appropriate LED on all of the physical disks in the expansion enclosure to blink.

To locate the expansion enclosure:

1. Select the **Hardware** tab.
2. Select a physical disk in the expansion enclosure that you want to locate.
3. Select **Hardware → Blink → Expansion Enclosure.**
   The LED or LEDs on the expansion enclosure or physical disks come on.
4. When you have located the expansion enclosure, click **OK**.
   The LEDs stop blinking. (If you have an expansion enclosure with a blue LED, the LED goes off). If any other blink operations (Blink Storage Array, Blink Disk Group, Blink Physical Disk Ports, Blink Expansion Enclosure, or Blink Physical Disk) are currently being invoked from another storage management station, these LEDs also stop blinking.
5. If the LEDs on the expansion enclosure do not stop blinking, from the AMW, select **Hardware → Blink → Stop All Indications.**
   If the LEDs successfully stop blinking, a confirmation message is displayed.
6. Click **OK**.

Capturing The State Information

Use the **Capture State Information** option to capture information about the current state of your storage array and save the captured information to a text file. You can then send the captured information to your Technical Support representative for analysis.

CAUTION: Potential to cause an unresponsive storage array – The Capture State option can cause a storage array to become unresponsive to both the host and the storage management station. Use this option only under the guidance of your Technical Support representative.

1. From the AMW, select **Monitor → Health → Capture State Information.**
2. Read the information in the **Confirm State Capture** dialog, and type `yes` to continue.
3. In the **Specify filename** text box, enter a name for the file to be saved, or browse to a previously saved file if you want to overwrite an existing file.

   Use the convention filename.dmp for the name of the file. The suffix .dmp is added to the file automatically if you do not specify a suffix for the file.
4. Click **Start**.

   NOTE: Each test shows a status of Executing while it is in progress. The test then shows Completed when it successfully finishes. If any of the tests cannot be completed, a Failed status is displayed in the Execution summary window.
5. Monitor the progress and completion status of all of the tests. When they finish, click **OK** to close the **State Capture** dialog.

   Clicking **Cancel** stops the state capture process, and any remaining tests do not complete. Any test information that has been generated to that point is saved to the state capture file.

   NOTE: See the online help topics for more information on troubleshooting, and recovering from failures.

SMrepassist Utility

SMrepassist (replication assistance) is a host-based utility for Windows platforms. This utility is installed with MD Storage Manager. Use this utility before and after you create a virtual disk copy on a Windows operating system to ensure that all the memory-resident data for file systems on the target virtual disk is flushed and that the driver recognizes signatures and file system partitions. You can also use this utility to resolve duplicate signature problems for snapshot virtual disks.

From a command prompt window on a host running Windows, navigate to: **C:\Program Files\Dell\MD Storage Manager \util** and run the following command:

```
SMrepassist -f <filesystem-identifier>
```

Where, `-f` flushes all the memory-resident data for the file system indicated by <filesystem-identifier>, and <filesystem-identifier>, specifies a unique file system in the following syntax: `drive-letter:<mount-point-path>`

The file system identifier may consist of only a drive letter, as in the following example:

```
SMrepassist -f E:
```

NOTE: In Windows, the mount point path is a drive letter.

An error message is displayed in the command line when the utility cannot distinguish between the following:

- Source virtual disk and snapshot virtual disk (for example, if the snapshot virtual disk has been removed).
- Standard virtual disk and virtual disk copy (for example, if the virtual disk copy has been removed).

Unidentified Devices

An unidentified node or device occurs when the MD Storage Manager cannot access a new storage array. Causes for this error include network connection problems, the storage array is turned off, or the storage array does not exist.

NOTE: Before beginning any recovery procedure, make sure that the Host Context Agent software is installed and running. If you started the host before the host was connected to the storage array, the Host Context Agent software is not able to find the storage array. If so, make sure that the connections are tight, and restart the Host Context Agent software.

- If a storage array is managed by using both out-of-band management and in-band management using the same host, a management network connection problem may prevent direct communication with the storage array. However, you may still be able to manage the storage array over the in-band connections. The opposite situation can also occur.
- If a storage array is managed through more than one host, it is possible that the storage array may become unresponsive to communication over the connections given by one host. However, you may still be able to manage the storage array over the connections provided by another host.

Recovering From An Unidentified Storage Array

To recover from an unidentified storage array:

1. Make sure that the network connection to the storage management station is functional.
2. Make sure that the controllers are installed and that the power to the storage array is turned on. Correct any existing problems before continuing.
3. If you have an in-band storage array, use the following procedure. Click **Refresh** after each step to check the results:
   a) Make sure that the Host Context Agent software is installed and running. If you started the host before the host was connected to the controllers in the storage array, the Host Context Agent software is not able to find the controllers. If so, make sure that the connections are tight, and restart the Host Context Agent software.
   b) Make sure that the network can access the host by using the `ping` command in the following syntax: `ping <host-name-or-IP-address-of-the-host>`

      If the network can access the host, continue to step c. If the network cannot access the host, skip to step d.
   c) Remove the host with the unresponsive status from the MD Storage Manager, and add that host again.

      If the host returns to optimal status, you have completed this procedure.
   d) Make sure that the power to the host is turned on and that the host is operational.
   e) If applicable, make sure that the host bus adapters have been installed in the host.
   f) Examine all external cables and switches or hubs to make sure that you cannot see any damage and that they are tightly connected.
   g) If you have recently replaced or added the controller, restart the Host Context Agent software so that the new controller is found.

      If a problem exists, make the appropriate modifications to the host.
4. If you have an out-of-band storage array, use the following procedure. Click **Refresh** after each step to make sure of the results:
   a) Make sure that the network can access the controllers by using the `ping` command. Use the following syntax: `ping <controller-IP-address>`

      If the network can access the controllers, continue to step b. If the network cannot access the controllers, skip to step c.
   b) Remove the storage array with the unresponsive status from MD Storage Manager, and add that storage array again.

      If the storage array returns to optimal status, you have completed this procedure.

c) Examine the Ethernet cables to make sure that you cannot see any damage and that they are tightly connected.
d) Make sure that the applicable network configuration tasks have been done (for example, the IP addresses have been assigned to each controller).

5. Make sure that the controller firmware is compatible with MD Storage Manager on your management station. If the controller firmware was upgraded, the MD Storage Manager may not have access to the storage array. A new version of MD Storage Manager may be needed to manage the storage array with the new version of the controller firmware.

   If this problem exists, see Getting Help.

6. Look to see if there is too much network traffic to one or more controllers. This problem corrects itself because the MD Storage Manager tries to re-establish communication with the controllers in the storage array at regular times. If the storage array was unresponsive and a subsequent attempt to connect to the storage array succeeds, the storage array becomes responsive.

7. For an out-of-band storage array, look to see if management operations are taking place on the storage array from other storage management stations. The type of management operations being done and the number of management sessions taking place together establish the number of TCP/IP connections made to a controller. When the maximum number of TCP/IP connections have been made, the controller stops responding. This problem corrects itself because after some TCP/IP connections are complete, the controller becomes responsive to other connection tries.

8. If the storage array is still unresponsive, problems may exist with the controllers.

   If these problems persist, see Getting Help.

Starting Or Restarting The Host Context Agent Software

The Host Context Agent software module is the software component that resides on the server or management station that communicates with the MD Series storage arrays. The SMagent software automatically starts after you reboot the host.

Starting The SMagent Software In Windows

1. Do one of the following:

   - Click **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**
   - Click **Start** → **Administrative Tools** → **Services**

2. In the **Services** dialog, select **Modular Disk Storage Manager Agent**.
3. If the modular disk storage manager agent is running, click **Action** → **Stop** and then wait approximately 5 seconds.
4. Click **Action** → **Start.**

Starting The SMagent Software In Linux

To start or restart the Host Context Agent software in Linux, enter the following command at the prompt:

```
SMagent start
```

The SMagent software may take a little time to initialize. The cursor is shown, but the terminal window does not respond. When the program starts, the following message is displayed:

```
SMagent started.
```

After the program completes the startup process, text similar to the following, is displayed:

```
Modular Disk Storage Manager Agent, Version 90.02.A6.14
```

```
Copyright (C) 2009-2010 Dell, Inc. All rights reserved.

Checking device <n/a> (/dev/sg10): Activating

Checking device /dev/sdb (/dev/sg11): Skipping

Checking device <n/a> (/dev/sg3): Activating

Checking device <n/a> (/dev/sg4): Activating

Checking device <n/a> (/dev/sg5): Activating

Checking device <n/a> (/dev/sg6): Activating

Checking device <n/a> (/dev/sg7): Activating

Checking device <n/a> (/dev/sg8): Activating

Checking device <n/a> (/dev/sg9): Activating
```

21

Getting Help

Contacting Dell

NOTE: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical support, or customer service issues:

1. Visit **support.dell.com**.
2. Select your support category.
3. If you are not a U.S. customer, select your country code at the bottom of the **support.dell.com** page, or select **All** to see more choices.
4. Select the appropriate service or support link based on your need.